REPORT DOCUMENTATION PAGE

Form Approved
OPM No. 0704-0188

AD-A225 460

1. AGENCY USE ONLY (Leave Blank) DTIC FILE COPY	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED Final 15 Jan 1990 to 15 Jan 1991

4. TITLE AND SUBTITLE Ada Compiler Validation Summary Report: Alsys, AlsyCOMP_040, Version 4.3, IBM PC/AT (Host) to TACCS AN/TYQ-33 (V)(Target), 900115A1.10241	5. FUNDING NUMBERS ②
6. AUTHOR(S) AFNOR, Paris, FRANCE	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) AFNOR Tour Europe, Cedex 7 F-92080 Paris la Defense FRANCE	8. PERFORMING ORGANIZATION REPORT NUMBER AVF-VSR-AFNOR-90-02
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Ada Joint Program Office United States Department of Defense Washington, D.C. 20301-3081	10. SPONSORING/MONITORING AGENCY REPORT NUMBER
11. SUPPLEMENTARY NOTES	
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.	12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

Alsys, AlsyCOMP_040, Version 4.3, Paris, France, IBM PC/AT under MS/DOS, Version 3.2 (Host) to TACCS AN/TYQ-33(V) under BTOS, Version 7.0 (Target), ACVC 1.10.

14. SUBJECT TERMS Ada programming language, Ada Compiler Validation Summary Report, Ada Compiler Validation Capability, Validation Testing, Ada Validation Office, Ada Validation Facility, ANSI/MIL-STD-1815A, Ada Joint Program Office			15. NUMBER OF PAGES
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT

AVF Control Number: AVF-VSR-AFNOR-90-02

Ada COMPILER VALIDATION SUMMARY REPORT:

Certificate Number: 900115A1.10241
Alsys
AlsyCOMP_040, Version 4.3
IBM PC/AT Host and TACCS AN/TYQ-33(V) Target

Completion of On-Site Testing:
15 January 1990

Prepared By:
AFNOR
Tour Europe
Cedex 7
F-92049 Paris la Défense

Prepared For:
Ada Joint Program Office
United States Department of Defense
Washington DC 20301-3081

Ada Compiler Validation Summary Report:

Compiler Name: AlsyCOMP_040, Version 4.3

Certificate Number: 900115A1.10241

Host: IBM PC/AT under MS/DOS, Version 3.2

Target: TACCS AN/TYQ-33(V) under BTOS, Version 7.0

Testing Completed 15 January 1990 Using ACVC 1.10

This report has been reviewed and is approved.

F de Labareyre

AFNOR
Fabrice Garnier de Labareyre
Tour Europe
Cedex 7
F-92049 Paris la Défense

[illegible]

Ada Validation Organization
Director, Computer & Software Engineering Division
Institute for Defense Analyses
Alexandria, VA 22311

for [illegible]

Ada Joint Program Office
Dr. John Solomond
Director
Department of Defense
Washington DC 20301

TABLE OF CONTENTS

CHAPTER 1

INTRODUCTION

This Validation Summary Report (VSR) describes the extent to which a specific Ada compiler conforms to the Ada Standard, ANSI/MIL-STD-1815A. This report explains all technical terms used within it and thoroughly reports the results of testing this compiler using the Ada Compiler Validation Capability (ACVC). An Ada compiler must be implemented according to the Ada Standard, and any implementation-dependent features must conform to the requirements of the Ada Standard. The Ada Standard must be implemented in its entirety, and nothing can be implemented that is not in the Standard.

Even though all validated Ada compilers conform to the Ada Standard, it must be understood that some differences do exist between implementations. The Ada Standard permits some implementation dependencies--for example, the maximum length of identifiers or the maximum values of integer types. Other differences between compilers result from the characteristics of particular operating systems, hardware, or implementation strategies. All the dependencies observed during the process of testing this compiler are given in this report.

The information in this report is derived from the test results produced during validation testing. The validation process includes submitting a suite of standardized tests, the ACVC, as inputs to an Ada compiler and evaluating the results. The purpose of validating is to ensure conformity of the compiler to the Ada Standard by testing that the compiler properly implements legal language constructs and that it identifies and rejects illegal language constructs. The testing also identifies behavior that is implementation dependent, but is permitted by the Ada Standard. Six classes of tests are used. These tests are designed to perform checks at compile time, at link time, and during execution.

1.1 PURPOSE OF THIS VALIDATION SUMMARY REPORT

This VSR documents the results of the validation testing performed on an Ada compiler. Testing was carried out for the following purposes:

- To attempt to identify any language constructs supported by the compiler that do not conform to the Ada Standard
- To attempt to identify any language constructs not supported by the compiler but required by the Ada Standard
- To determine that the implementation-dependent behavior is allowed by the Ada Standard

Testing of this compiler was conducted by Alsys under the direction of the AVF according to procedures established by the Ada Joint Program Office and administered by the Ada Validation Organization (AVO). On-site testing was completed 15 January 1990 at Alsys Inc, in Burlington MA, USA.

1.2 USE OF THIS VALIDATION SUMMARY REPORT

Consistent with the national laws of the originating country, the AVO may make full and free public disclosure of this report. In the United States, this is provided in accordance with the "Freedom of Information Act" (5 U.S.C. #552). The results of this validation apply only to the computers, operating systems, and compiler versions identified in this report.

The organizations represented on the signature page of this report do not represent or warrant that all statements set forth in this report are accurate and complete, or that the subject compiler has no nonconformities to the Ada Standard other than those presented. Copies of this report are available to the public from:

Ada Information Clearinghouse
Ada Joint Program Office
OUSDRE
The Pentagon, Rm 3D-139 (Fern Street)
Washington DC 20301-3081

or from:

AFNOR
Tour Europe
cedex 7
F-92049 Paris la Défense

Questions regarding this report or the validation test results should be directed to the AVF listed above or to:

Ada Validation Organization
Institute for Defense Analyses
1801 North Beauregard Street
Alexandria VA 22311

1.3 REFERENCES

1. Reference Manual for the Ada Programming Language, ANSI/MIL-STD-1815A, February 1983, and ISO 8652-1987.

2. Ada Compiler Validation Procedures, Ada Joint Program Office, May 1989.

3. Ada Compiler Validation Capability Implementers' Guide, SofTech, Inc., December 1986.

4. Ada Compiler Validation Capability User's Guide, January 1989

1.4 DEFINITION OF TERMS

ACVC	The Ada Compiler Validation Capability. The set of Ada programs that tests the conformity of an Ada compiler to the Ada programming language.
Ada Commentary	An Ada Commentary contains all information relevant to the point addressed by a comment on the Ada Standard. These comments are given a unique identification number having the form AI-ddddd.
Ada Standard	ANSI/MIL-STD-1815A, February 1983 and ISO 8652-1987.
Applicant	The agency requesting validation.
AVF	The Ada Validation Facility. The AVF is responsible for conducting compiler validations according to procedures contained in the Ada Compiler Validation Procedures.
AVO	The Ada Validation Organization. The AVO has oversight authority over all AVF practices for the purpose of maintaining a uniform process for validation of Ada compilers. The AVO provides administrative and technical support for Ada validations to ensure consistent practices.
Compiler	A processor for the Ada language. In the context of this report, a compiler is any language processor, including cross-compilers, translators, and interpreters.
Failed test	An ACVC test for which the compiler generates a result that demonstrates nonconformity to the Ada Standard.
Host	The computer on which the compiler resides.
Inapplicable test	An ACVC test that uses features of the language that a compiler is not required to support or may legitimately support in a way other than the one expected by the test.
Passed test	An ACVC test for which a compiler generates the expected result.
Target	The computer which executes the code generated by the compiler.

Test	A program that checks a compiler's conformity regarding a particular feature or a combination of features to the Ada Standard. In the context of this report, the term is used to designate a single test, which may comprise one or more files.
Withdrawn	An ACVC test found to be incorrect and not used to check test conformity to the Ada Standard. A test may be incorrect because it has an invalid test objective, fails to meet its test objective, or contains illegal or erroneous use of the language.

1.5 ACVC TEST CLASSES

Conformity to the Ada Standard is measured using the ACVC. The ACVC contains both legal and illegal Ada programs structured into six test classes: A, B, C, D, E, and L. The first letter of a test name identifies the class to which it belongs. Class A, C, D, and E tests are executable, and special program units are used to report their results during execution. Class B tests are expected to produce compilation errors. Class L tests are expected to produce errors because of the way in which a program library is used at link time.

Class A tests ensure the successful compilation and execution of legal Ada programs with certain language constructs which cannot be verified at run time. There are no explicit program components in a Class A test to check semantics. For example, a Class A test checks that reserved words of another language (other than those already reserved in the Ada language) are not treated as reserved words by an Ada compiler. A Class A test is passed if no errors are detected at compile time and the program executes to produce a PASSED message.

Class B tests check that a compiler detects illegal language usage. Class B tests are not executable. Each test in this class is compiled and the resulting compilation listing is examined to verify that every syntax or semantic error in the test is detected. A Class B test is passed if every illegal construct that it contains is detected by the compiler.

Class C tests check the run time system to ensure that legal Ada programs can be correctly compiled and executed. Each Class C test is self-checking and produces a PASSED, FAILED, or NOT APPLICABLE message inc[illegible]g the result when it is executed.

Class D tests check the compilation and execution capacities of a compiler. Since there are no capacity requirements placed on a compiler by the Ada Standard for some parameters--for example, the number of identifiers permitted in a compilation or the number of units in a library--a compiler may refuse to compile a Class D test and still be a conforming compiler. Therefore, if a Class D test fails to compile because the capacity of the compiler is exceeded, the test is classified as inapplicable. If a Class D test compiles successfully, it is self-checking and produces a PASSED or FAILED message during execution.

Class E tests are expected to execute successfully and check implementation-dependent options and resolutions of ambiguities in the Ada Standard. Each Class E test is self-checking and produces a NOT APPLICABLE, PASSED, or FAILED message when it is compiled and executed. However, the Ada Standard permits an implementation to reject programs containing some features addressed by Class E tests during compilation. Therefore, a Class E test is passed by a compiler if it is compiled successfully and executes to produce a PASSED message, or if it is rejected by the compiler for an allowable reason.

Class L tests check that incomplete or illegal Ada programs involving multiple, separately compiled units are detected and not allowed to execute. Class L tests are compiled separately and execution is attempted. A Class L test passes if it is rejected at link time--that is, an attempt to execute the main program must generate an error message before any declarations in the main program or any units referenced by the main program are elaborated. In some cases, an implementation may legitimately detect errors during compilation of the test.

Two library units, the package REPORT and the procedure CHECK_FILE, support the self-checking features of the executable tests. The package REPORT provides the mechanism by which executable tests report PASSED, FAILED, or NOT APPLICABLE results. It also provides a set of identity functions used to defeat some compiler optimizations allowed by the Ada Standard that would circumvent a test objective. The procedure CHECK_FILE is used to check the contents of text files written by some of the Class C tests for Chapter 14 of the Ada Standard. The operation of REPORT and CHECK_FILE is checked by a set of executable tests. These tests produce messages that are examined to verify that the units are operating correctly. If these units are not operating correctly, then the validation is not attempted.

The text of each test in the ACVC follows conventions that are intended to ensure that the tests are reasonably portable without modification. For example, the tests make use of only the basic set of 55 characters, contain lines with a maximum length of 72 characters, use small numeric values, and place features that may not be supported by all implementations in separate tests. However, some tests contain values that require the test to be customized according to implementation-specific values--for example, an illegal file name. A list of the values used for this validation is provided in Appendix C.

A compiler must correctly process each of the tests in the suite and demonstrate conformity to the Ada Standard by either meeting the pass criteria given for the test or by showing that the test is inapplicable to the implementation. The applicability of a test to an implementation is considered each time the implementation is validated. A test that is inapplicable for one validation is not necessarily inapplicable for a subsequent validation. Any test that was determined to contain an illegal language construct or an erroneous language construct is withdrawn from the ACVC and, therefore, is not used in testing a compiler. The tests withdrawn at the time of this validation are given in Appendix D.

CHAPTER 2

CONFIGURATION INFORMATION

2.1 CONFIGURATION TESTED

The candidate compilation system for this validation was tested under the following configuration:

Compiler: AlsyCOMP_040, Version 4.3

ACVC Version: 1.10

Certificate Number: 900115A1.10241

Host Computer:

Machine:	IBM PC/AT
Operating System:	MS/DOS, Version 3.2
Memory Size:	640 Kb of main memory plus 5 Mb on extent memory

Target Computer:

Machine:	TACCS AN/TYQ-33(V)
Operating System:	BTOS, Version 7.0
Memory Size:	1024 Kb of main memory plus 5 Mb on extent memory

Communication Network:	Unisys BTOS Clustershare, Level 1.1

2.2 IMPLEMENTATION CHARACTERISTICS

One of the purposes of validating compilers is to determine the behavior of a compiler in those areas of the Ada Standard that permit implementations to differ. Class D and E tests specifically check for such implementation differences. However, tests in other classes also characterize an implementation. The tests demonstrate the following characteristics:

a. Capacities.

(1) The compiler correctly processes a compilation containing 723 variables in the same declarative part. (See test D29002K.)

(2) The compiler correctly processes tests containing loop statements nested to 65 levels. (See tests D55A03A..H (8 tests).)

(3) The compiler correctly processes a test containing block statements nested to 65 levels. (See test D56001B.)

(4) The compiler correctly processes tests containing recursive procedures separately compiled as subunits nested to 17 levels. (See tests D64005E..G (3 tests).)

b. Predefined types.

(1) This implementation supports the additional predefined types, SHORT_INTEGER, LONG_INTEGER, LONG_FLOAT in the package STANDARD. (See tests B86001T..Z (7 tests).)

c. Based literals.

(1) An implementation is allowed to raise NUMERIC_ERROR or CONSTRAINT_ERROR when a value exceeds SYSTEM.MAX_INT . This implementation raises CONSTRAINT_ERROR during execution. (See test E24201A.)

d. Expression evaluation.

The order in which expressions are evaluated and the time at which constraints are checked are not defined by the language. While the ACVC tests do not specifically attempt to determine the order of evaluation of expressions, test results indicate the following:

(1) Apparently no default initialization expressions for record components are evaluated before any value is checked to belong to a component's subtype. (See test C32117A.)

(2) Assignments for subtypes are performed with the same precision as the base type. (See test C35712B.)

(3) This implementation uses no extra bits for extra precision. This implementation uses all extra bits for extra range. (See test C35903A.)

(4) NUMERIC_ERROR is raised when an integer literal operand in a comparison or membership test is outside the range of the base type. (See test C45232A.)

(5) NUMERIC_ERROR is raised when a literal operand in a fixed-point comparison or membership test is outside the range of the base type. (See test C45252A.)

(6) Underflow is gradual. (See tests C45524A..Z.) (26 tests)

e. Rounding.

The method by which values are rounded in type conversions is not defined by the language. While the ACVC tests do not specifically attempt to determine the method of rounding, the test results indicate the following:

(1) The method used for rounding to integer is round to even. (See tests C46012A..Z.) (26 tests)

(2) The method used for rounding to longest integer is round to even. (See tests C46012A..Z.) (26 tests)

(3) The method used for rounding to integer in static universal real expressions is round to even. (See test C4A014A.)

f. Array types.

An implementation is allowed to raise NUMERIC_ERROR or CONSTRAINT_ERROR for an array having a 'LENGTH that exceeds STANDARD.INTEGER'LAST and/or SYSTEM.MAX_INT. For this implementation:

(1) Declaration of an array type or subtype declaration with more than SYSTEM.MAX_INT components raises NUMERIC_ERROR sometimes, CONSTRAINT_ERROR sometimes. (See test C36003A.)

(2) CONSTRAINT_ERROR is raised when 'LENGTH is applied to an array type with INTEGER'LAST + 2 components. (See test C36202A.)

(3) CONSTRAINT_ERROR is raised when an array type with SYSTEM.MAX_INT + 2 components is declared. (See test C36202B.)

(4) A packed BOOLEAN array having a 'LENGTH exceeding INTEGER'LAST raises no exception. (See test C52103X.)

(5) A packed two-dimensional BOOLEAN array with more than INTEGER'LAST components raises CONSTRAINT_ERROR when the length of a dimension is calculated and exceeds INTEGER'LAST. (See test C52104Y.)

(6) In assigning one-dimensional array types, the expression is evaluated in its entirety before CONSTRAINT_ERROR is raised when checking whether the expression's subtype is compatible with the target's subtype. (See test C52013A.)

(7) In assigning two-dimensional array types, the expression is not evaluated in its entirety before CONSTRAINT_ERROR is raised when checking whether the expression's subtype is compatible with the target's subtype. (See test C52013A.)

g. A null array with one dimension of length greater than INTEGER'LAST may raise NUMERIC_ERROR or CONSTRAINT_ERROR either when declared or assigned. Alternatively, an implementation may accept the declaration. However, lengths must match in array slice assignments. This implementation raises no exception. (See test E52103Y.)

h. Discriminated types.

(1) In assigning record types with discriminants, the expression is evaluated in its entirety before CONSTRAINT_ERROR is raised when checking whether the expression's subtype is compatible with the target's subtype. (See test C52013A.)

i. Aggregates.

(1) In the evaluation of a multi-dimensional aggregate, all choices appear to be evaluated before checking against the index type. (See tests C43207A and C43207B.)

(2) In the evaluation of an aggregate containing subaggregates, not all choices are evaluated before being checked for identical bounds. (See test E43212B.)

(3) CONSTRAINT_ERROR is raised after all choices are evaluated when a bound in a non-null range of a non-null aggregate does not belong to an index subtype. (See test E43211B.)

j. Pragmas.

(1) The pragma INLINE is supported for functions or procedures, but not functions called inside a package specification. (See tests LA3004A..B, EA3004C..D, and CA3004E..F.)

k. Generics.

(1) Generic specifications and bodies can be compiled in separate compilations. (See tests CA1012A, CA2009C, CA2009F, BC3204C, and BC3205D.)

(2) Generic subprogram declarations and bodies can be compiled in separate compilations. (See tests CA1012A and CA2009F.)

(3) Generic library subprogram specifications and bodies can be compiled in separate compilations. (See test CA1012A.)

(4) Generic non-library package bodies as subunits can be compiled in -arate compilations. (See test CA2009C.)

(5) Generic non-library subprogram bodies can be compiled in separate compilations from their stubs. (See test CA2009F.)

(6) Generic unit bodies and their subunits can be compiled in separate compilations. (See test CA3011A.)

(7) Generic package declarations and bodies can be compiled in separate compilations. (See tests CA2009C, BC3204C, and BC3205D.)

(8) Generic library package specifications and bodies can be compiled in separate compilations. (See tests BC3204C and BC3205D.)

(9) Generic unit bodies and their subunits can be compiled in separate compilations. (See test CA3011A.)

l. Input and output.

(1) The package SEQUENTIAL_IO can be instantiated with unconstrained array types and record types with discriminants without defaults. (See tests AE2101C, EE2201D, and EE2201E.)

(2) The package DIRECT_IO can be instantiated with unconstrained array types and record types with discriminants but CREATE will raise USE_ERROR. (See tests AE2101H, EE2401D, and EE2401G.)

(4) Modes IN_FILE and OUT_FILE are supported for SEQUENTIAL_IO but not CREATE in mode IN_FILE. (See tests CE2102D..E, CE2102N, and CE2102P.)

(5) Modes IN_FILE, OUT_FILE, and INOUT_FILE are supported for DIRECT_IO but not CREATE in mode IN_FILE. (See tests CE2102F, CE2102I..J, CE2102R, CE2102T, and CE2102V.)

(6) Modes IN_FILE and OUT_FILE are supported for text files but not CREATE in mode IN_FILE. (See tests CE3102E and CE3102I..K.)

(7) RESET from OUT_FILE to IN_FILE only and DELETE operations are supported for SEQUENTIAL_IO. (See tests CE2102G and CE2102X.)

(8) RESET except from IN_FILE to INOUT_FILE or to OUT_FILE and DELETE operations are supported for DIRECT_IO. (See tests CE2102K and CE2102Y.)

(9) RESET and DELETE operations are supported for text files. (See tests CE3102F..G, CE3104C, CE3110A, and CE3114A.)

(10) Overwriting to a sequential file truncates to the last element written. (See test CE2208B.)

(11) Temporary sequential files are given names and not deleted when closed. (See test CE2108A.)

(12) Temporary direct files are given names and not deleted when closed. (See test CE2108C.)

(13) Temporary text files are given names and not deleted when closed. (See test CE3112A.)

(14) More than one internal file can be associated with each external file for sequential files when reading only. (See tests CE2107A..E, CE2102L, CE2110B, and CE2111D.)

(15) More than one internal file can be associated with each external file for direct files when reading only. (See tests CE2107F..H, CE2110D and CE2111H.)

(16) More than one internal file can be associated with each external file for text files when reading only. (See tests CE3111A..E, CE3114B, and CE3115A.)

CHAPTER 3

TEST INFORMATION

3.1 TEST RESULTS

Version 1.10 of the ACVC comprises 3717 tests. When this compiler was tested, 44 tests had been withdrawn because of test errors. The AVF determined that 368 tests were inapplicable to this implementation. All inapplicable tests were processed during validation testing except for 201 executable tests that use floating-point precision exceeding that supported by the implementation. Modifications to the code, processing, or grading for 53 tests were required. (See section 3.6.)

The AVF concludes that the testing results demonstrate acceptable conformity to the Ada Standard.

3.2 SUMMARY OF TEST RESULTS BY CLASS

RESULT	TEST CLASS						TOTAL
	A	B	C	D	E	L	
Passed	129	1132	1955	17	26	46	3305
Inapplicable	0	6	360	0	2	0	368
Withdrawn	1	2	35	0	6	0	44
TOTAL	130	1140	2350	17	34	46	3717

3.3 SUMMARY OF TEST RESULTS BY CHAPTER

RESULT	CHAPTER													TOTAL
	2	3	4	5	6	7	8	9	10	11	12	13	14	
Passed	198	577	555	248	172	99	161	331	137	36	252	259	280	3305
Inappl	14	72	125	0	0	0	5	1	0	0	0	110	41	368
Wdrn	1	1	0	0	0	0	0	2	0	0	1	35	4	44
TOTAL	213	650	680	248	172	99	166	334	137	36	253	404	325	3717

3.4 WITHDRAWN TESTS

The following 44 tests were withdrawn from ACVC Version 1.10 at the time of this validation:

A39005G	B97102E	BC3009B	C97116A	CD2A62D	CD2A63A	CD2A63B	CD2A63C	CD2A63D
CD2A66A	CD2A66B	CD2A66C	CD2A66D	CD2A73A	CD2A73B	CD2A73C	CD2A73D	CD2A76A
CD2A76B	CD2A76C	CD2A76D	CD2A81G	CD2A83G	CD2A84M	CD2A84N	CD2D11B	CD2B15C
CD5007B	CD50110	CD7105A	CD7203B	CD7204B	CD7205C	CD7205D	CE2107I	CE3111C
CE3301A	CE3411B	E28005C	ED7004B	ED7005C	ED7005D	ED7006C	ED7006D	

See Appendix D for the reason that each of these tests was withdrawn.

3.5 INAPPLICABLE TESTS

Some tests do not apply to all compilers because they make use of features that a compiler is not required by the Ada Standard to support. Others may depend on the result of another test that is either inapplicable or withdrawn. The applicability of a test to an implementation is considered each time a validation is attempted. A test that is inapplicable for one validation attempt is not necessarily inapplicable for a subsequent attempt. For this validation attempt, 368 tests were inapplicable for the reasons indicated:

- The following 201 tests are not applicable because they have floating-point type declarations requiring more digits than System.Max_Digits:

C24113L..Y (14 tests)	C35705L..Y (14 tests)
C35706L..Y (14 tests)	C35707L..Y (14 tests)
C35708L..Y (14 tests)	C35802L..Z (15 tests)
C45241L..Y (14 tests)	C45321L..Y (14 tests)
C45421L..Y (14 tests)	C45521L..Z (15 tests)
C45524L..Z (15 tests)	C45621L..Z (15 tests)
C45641L..Y (14 tests)	C46012L..Z (15 tests)

- C35702A and B86001T are not applicable because this implementation supports no predefined type Short_Float.

- C45531M..P (4 tests) and C45532M..P (4 tests) are not applicable because the value of System.Max_Mantissa is less than 32.

- C86001F, is not applicable because recompilation of Package SYSTEM is not allowed.

- B86001X, C45231D, and CD7101G are not applicable because this implementation does not support any predefined integer type with a name other than Integer, Long_Integer, or Short_Integer.

- B86001Y is not applicable because this implementation supports no predefined fixed-point type other than Duration.

- B86001Z is not applicable because this implementation supports no predefined floating-point type with a name other than Float, Long_Float, or Short_Float.

- B91001H is not applicable because address clause for entries is not supported by this implementation.

- BD5006D is not applicable because address clause for packages is not supported by this implementation.

- The following 10 tests are not applicable because size clause on float is not supported by this implementation:

CD1009C	CD2A41A..B (2 tests)
CD2A41E	CD2A42A..B (2 tests)
CD2A42E..F (2 tests)	CD2A42I..J (2 tests)

- CD1C04B, CD1C04E, CD4051A..D (4 tests) are not applicable because representation clause on derived records or derived tasks is not supported by this implementation.

- CD2A84B..I (8 tests), CD2A84K..L (2 tests) are not applicable because size clause on access type is not supported by this implementation.

- The following 28 tests are not applicable because size clause for derived private type is not supported by this implementation:

CD1C04A	CD2A21C..D (2 tests)
CD2A22C..D (2 tests)	CD2A22G..H (2 tests)
CD2A31C..D (2 tests)	CD2A32C..D (2 tests)
CD2A32G..H (2 tests)	CD2A41C..D (2 tests)
CD2A42C..D (2 tests)	CD2A42G..H (2 tests)
CD2A51C..D (2 tests)	CD2A52C..D (2 tests)
CD2A52G..H (2 tests)	CD2A53D
CD2A54D	CD2A54H

- The following 29 tests are not applicable because of the way this implementation allocates storage space for one component, size specification clause for an array type or for a record type requires compression of the storage space needed for all the components (without gaps).

CD2A61A..D (4 tests)	CD2A61F
CD2A61H..L (5 tests)	CD2A62A..C (3 tests)
CD2A71A..D (4 tests)	CD2A72A..D (4 tests)
CD2A74A..D (4 tests)	CD2A75A..D (4 tests)

- CD4041A is not applicable because alignment "at mod 8" is not supported by this implementation.

- The following 21 tests are not applicable because address clause for a constant is not supported by this implementation:

CD5011B,D,F,H,L,N,R (7 tests)	CD5012C,D,G,H,L (5 tests)
CD5013B,D,F,H,L,N,R (7 tests)	CD5014U,W (2 tests)

- CD5012J, CD5013S, CD5014S are not applicable because address clause for a task is not supported by this implementation.

- CE2102E is inapplicable because this implementation supports create with out_file mode for SEQUENTIAL_IO.

- CE2102F is inapplicable because this implementation supports create with inout_file mode for DIRECT_IO.

- CE2102J is inapplicable because this implementation supports create with out_file mode for DIRECT_IO.
- CE2102N is inapplicable because this implementation supports open with in_file mode for SEQUENTIAL_IO.
- CE2102O is inapplicable because this implementation supports RESET with in_file mode for SEQUENTIAL_IO.
- CE2102P is inapplicable because this implementation supports open with out_file mode for SEQUENTIAL_IO.
- CE2102Q is inapplicable because this implementation supports RESET with out_file mode for SEQUENTIAL_IO.
- CE2102R is inapplicable because this implementation supports open with inout_file mode for DIRECT_IO.
- CE2102S is inapplicable because this implementation supports RESET with inout_file mode for DIRECT_IO.
- CE2102T is inapplicable because this implementation supports open with in_file mode for DIRECT_IO.
- CE2102U is inapplicable because this implementation supports RESET with in_file mode for DIRECT_IO.
- CE2102V is inapplicable because this implementation supports open with out_file mode for DIRECT_IO.
- CE2102W is inapplicable because this implementation supports RESET with out_file mode for DIRECT_IO.
- CE2105A is inapplicable because CREATE with IN_FILE mode is not supported by this implementation for SEQUENTIAL_IO.
- CE2105B is inapplicable because CREATE with IN_FILE mode is not supported by this implementation for DIRECT_IO.
- The following 6 tests are not applicable because multiple internal files cannot be associated with the same external file when one or more files is writing for sequential files. The proper exception is raised when multiple access is attempted.

  ```
  CE2107B..E (4 tests)    CE2107L
  CE2110B
  ```
- The following 4 tests are not applicable because multiple internal files cannot be associated with the same external file when one or more files is writing for direct files. The proper exception is raised when multiple access is attempted.

  ```
  CE2107G..H (2 tests)    CE2110D
  CE2111H
  ```
- CE2111C..D (2 tests) are not applicable because reseting from IN_FILE to OUT_FILE mode for sequential files is not supported.
- EE2401D and EE2401G are not applicable because USE_ERROR is raised when the CREATE of an instantiation of DIRECT_IO with unconstrained type is called.

- CE2401H is not applicable because create with inout_file mode for unconstrained records with default discriminants is not supported by this implementation.

- CE3102F is inapplicable because this implementation supports reset for text files, for out_file, in_file and from out_file to in_file mode.

- CE3102G is inapplicable because this implementation supports deletion of an external file for text files.

- CE3102I is inapplicable because this implementation supports create with out_file mode for text files.

- CE3102J is inapplicable because this implementation supports open with in_file mode for text files.

- CE3102K is inapplicable because this implementation supports open with out_file mode for text files.

- CE3109A is inapplicable because text file CREATE with IN_FILE mode is not supported by this implementation.

- The following 5 tests are not applicable because multiple internal files cannot be associated with the same external file when one or more files is writing for text files. The proper exception is raised when multiple access is attempted.

CE3111B	CE3111D..E (2 tests)
CE3114B	CE3115A

3.6 TEST, PROCESSING, AND EVALUATION MODIFICATIONS

It is expected that some tests will require modifications of code, processing, or evaluation in order to compensate for legitimate implementation behavior. Modifications are made by the AVF in cases where legitimate implementation behavior prevents the successful completion of an (otherwise) applicable test. Examples of such modifications include: adding a length clause to alter the default size of a collection; splitting a Class B test into subtests so that all errors are detected; and confirming that messages produced by an executable test demonstrate conforming behavior that wasn't anticipated by the test (such as raising one exception instead of another).

Modifications were required for 53 tests.

The following 30 tests were split because syntax errors at one point resulted in the compiler not detecting other errors in the test:
B23004A B24007A B24009A B25002A B26005A B27005A B28003A B32202A B32202B B32202C
B33001A B36307A B37004A B49003A B49005A B61012A B62001B B74304B B74304C B74401F
B74401R B91004A B95032A B95069A B95069B BA1101B BC2001D BC3009A BC3009C BD5005B

The following 21 tests were split in order to show that the compiler was able to find the representation clause indicated by the comment --N/A =>ERROR :
CD2A61A CD2A61B CD2A61F CD2A61I CD2A61J CD2A62A CD2A62B CD2A71A CD2A71B CD2A72A
CD2A72B CD2A75A CD2A75B CD2A84B CD2A84C CD2A84D CD2A84E CD2A84F CD2A84G CD2A84H
CD2A84I

The test EA3004D when run as it is, the implementation fails to detect an error on line 27 of test file EA3004D6M . This is because the pragma INLINE has no effect when its object is within a package specification. However, the results of running the test as it is does not confirm that the pragma had no effect, only that the package was not made obsolete. By re-ordering the compilations so that the two subprograms are compiled after file D5 (the re-compilation of the "with"ed package that makes the various earlier units obsolete), we create a test that shows that indeed pragma INLINE has no effect when applied to a subprogram that is called within a package specification: the test then executes and produces the expected NOT_APPLICABLE result (as though INLINE were not supported at all). The re-ordering of EA3004D test files is 0-1-4-5-2-3-6.

BA2001E requires that duplicate names of subunits with a common ancestor be detected and rejected at compile time. This implementation detects the error at link time, and the AVO ruled that this behavior is acceptable.

3.7 ADDITIONAL TESTING INFORMATION

3.7.1 Prevalidation

Prior to validation, a set of test results for ACVC Version 1.10 produced by the AlsyCOMP_040, Version 4.3 compiler was submitted to the AVF by the applicant for review. The full set of test results produced by the compiler was compared with a set of test results from a validated compiler. Analysis of the comparison results demonstrated that the compiler successfully passed all applicable tests, and the compiler exhibited the expected behavior on all inapplicable tests.

3.7.2 Test Method

Testing of the AlsyCOMP_040, Version 4.3 compiler using ACVC Version 1.10 was conducted on-site by a validation team from the AVF. The configuration in which the testing was performed is described by the following designations of hardware and software components:

Host computer:	IBM PC/AT
Host operating system:	MS/DOS, Version 3.2
Target computer:	TACCS AN/TYQ-33(V)
Target operating system:	BTOS, Version 7.0
Compiler:	AlsyCOMP_040, Version 4.3

The host and target computers were linked via Unisys BTOS Clustershare, Level 1.1.

The full set of tests for ACVC Version 1.10 except withdrawn tests and tests requiring unsupported floating-point precisions (tests that make use of implementation-specific values were customized before) was compiled by the AlsyCOMP_040, Version 4.3 and linked on the IBM PC/AT, then all executable images were transferred to the TACCS AN/TYQ-33(V) via Unisys BTOS Clustershare, Level 1.1. and run.

The full set of test results was compared on the host with the full set of prevalidation test results. No difference occurred.

The compiler was tested using command scripts provided by Alsys and reviewed by the validation team. The compiler was tested using all default option settings except for the following:

OPTION	EFFECT
CALLS=INLINED	Allow inline insertion of code for subprograms and take pragma INLINE into account
GENERIC=STUBS	Code of generic instantiation is placed in separate units.

Tests were compiled, linked, and executed (as appropriate) using two host computers and one target computer. Test output, compilation listings, and job logs were captured on diskettes and archived at the AVF. The listings examined on-site by the validation team were also archived.

3.7.3 Test Site

Testing was conducted at Alsys Inc, in Burlington MA, USA and was completed on 15 January 1990.

APPENDIX A

DECLARATION OF CONFORMANCE

Alsys has submitted the following Declaration of Conformance concerning the AlsyCOMP_040, Version 4.3 compiler.

DECLARATION OF CONFORMANCE

Compiler Implementor:	Alsys
Ada Validation Facility:	AFNOR, Tour Europe Cedex 7, F-92080 Paris la Défense

Ada Compiler Validation Capability (ACVC) Version: 1.10

Base Configuration

Base Compiler Name:	AlsyCOMP_040, Version 4.3
Host Architecture:	IBM AT
Host OS and Version:	MS/DOS, Version 3.2
Target Architecture:	TACCS AN/TYQ-33(V)
Target OS and Version:	BTOS, Version 7.0

Implementor's Declaration

I, the undersigned, representing Alsys, have implemented no deliberate extensions to the Ada Language Standard ANSI/MIL-STD-1815A in the compiler(s) listed in this declaration. I declare that Alsys is the owner of record of the Ada language compiler(s) listed above and, as such, is responsible for maintaining said compiler(s) in conformance to ANSI/MIL-STD-1815A. All certificates and registrations for Ada language compiler(s) listed in this declaration shall be made only in the owner's corporate name.

[illegible] Date: [illegible]

Alsys
Mike Blanchette, Vice President and Director of Engineering

Owner's Declaration

I, the undersigned, representing Alsys, take full responsibility for implementation and maintenance of the Ada compiler(s) listed above, and agree to the public disclosure of the final Validation Summary Report. I declare that all of the Ada language compilers listed, and their host/target performance, are in compliance with the Ada Language Standard ANSI/MIL-STD-1815A.

[illegible] Date: [illegible]

Alsys
Mike Blanchette, Vice President and Director of Engineering

APPENDIX B

TEST PARAMETERS

Certain tests in the ACVC make use of implementation-dependent values, such as the maximum length of an input line and invalid file names. A test that makes use of such values is identified by the extension .TST in its file name. Actual values to be substituted are represented by names that begin with a dollar sign. A value must be substituted for each of these names before the test is run. The values used for this validation are given below.

Name and Meaning	Value
$ACC_SIZE An integer literal whose value is the number of bits sufficient to hold any value of an access type.	32
$BIG_ID1 Identifier the size of the maximum input line length with varying last character.	(254 * 'A') & '1'
$BIG_ID2 Identifier the size of the maximum input line length with varying last character.	(254 * 'A) & '2'
$BIG_ID3 Identifier the size of the maximum input line length with varying middle character.	(126 * 'A') & '3' & (128 * 'A')
$BIG_ID4 Identifier the size of the maximum input line length with varying middle character.	(126 * 'A') & '4' & (128 * 'A')

Name and Meaning	Value
$BIG_INT_LIT An integer literal of value 298 with enough leading zeroes so that it is the size of the maximum line length.	(252 * '0') & '298'
$BIG_REAL_LIT A universal real literal of value 690.0 with enough leading zeroes to be the size of the maximum line length.	(250 * '0') & '690.0'
$BIG_STRING1 A string literal which when catenated with BIG_STRING2 yields the image of BIG_ID1.	'"' & (127 * 'A') & '"'
[illegible]ING2 A string literal which when catenated to the end of BIG_STRING1 yields the image of BIG_ID1.	'"' & (127 * 'A') & '1"'
$BLANKS A sequence of blanks twenty characters less than the size of the maximum line length.	(235 * ' ')
$COUNT_LAST A universal integer literal whose value is TEXT_IO.COUNT'LAST.	2147483647
$DEFAULT_MEM_SIZE An integer literal whose value is SYSTEM.MEMORY_SIZE.	655360
$DEFAULT_STOR_UNIT An integer literal whose value is SYSTEM.STORAGE_UNIT.	8
$DEFAULT_SYS_NAME The value of the constant SYSTEM.SYSTEM_NAME.	I_80X86
$DELTA_DOC A real literal whose value is SYSTEM.FINE_DELTA.	2#1.0#E-31

Name and Meaning	Value
$FIELD_LAST A universal integer literal whose value is TEXT_IO.FIELD'LAST.	255
$FIXED_NAME The name of a predefined fixed-point type other than DURATION.	NO_SUCH_FIXED_TYPE
$FLOAT_NAME The name of a predefined floating-point type other than FLOAT, SHORT_FLOAT, or LONG_FLOAT.	NO_SUCH_TYPE
$GREATER_THAN_DURATION A universal real literal that lies between DURATION'BASE'LAST and DURATION'LAST or any value in the range of DURATION.	2_097_151.999_023_437_51
$GREATER_THAN_DURATION_BASE_LAST A universal real literal that is greater than DURATION'BASE'LAST.	3_000_000.0
$HIGH_PRIORITY An integer literal whose value is the upper bound of the range for the subtype SYSTEM.PRIORITY.	10
$ILLEGAL_EXTERNAL_FILE_NAME1 An external file name specifying a non existent directory	ILLEGAL\!@#$%^&*()/_+~
$ILLEGAL_EXTERNAL_FILE_NAME2 An external file name different from $ILLEGAL_EXTERNAL_FILE_NAME1	!@#$%^&*()?/)(*&^\!@#$%^
$INTEGER_FIRST A universal integer literal whose value is INTEGER'FIRST.	-32768
$INTEGER_LAST A universal integer literal whose value is INTEGER'LAST.	32767
$INTEGER_LAST_PLUS_1 A universal integer literal whose value is INTEGER'LAST + 1.	32768

Name and Meaning	Value
$LESS_THAN_DURATION A universal real literal that lies between DURATION'BASE'FIRST and DURATION'FIRST or any value in the range of DURATION.	-2_097_152.5
$LESS_THAN_DURATION_BASE_FIRST A universal real literal that is less than DURATION'BASE'FIRST.	-3_000_000.0
$LOW_PRIORITY An integer literal whose value is the lower bound of the range for the subtype SYSTEM.PRIORITY.	1
$MANTISSA_DOC An integer literal whose value is SYSTEM.MAX_MANTISSA.	31
$MAX_DIGITS Maximum digits supported for floating-point types.	15
$MAX_IN_LEN Maximum input line length permitted by the implementation.	255
$MAX_INT A universal integer literal whose value is SYSTEM.MAX_INT.	2147483647
$MAX_INT_PLUS_1 A universal integer literal whose value is SYSTEM.MAX_INT+1.	2147483648
$MAX_LEN_INT_BASED_LITERAL A universal integer based literal whose value is 2:11: with enough leading zeroes in the mantissa to be MAX_IN_LEN long.	'2:' & (250 * '0') & '11:'
$MAX_LEN_REAL_BASED_LITERAL A universal real based literal whose value is 16: F.E: with enough leading zeroes in the mantissa to be MAX_IN_LEN long.	'16:' & (248 * '0') & 'F.E:'

Name and Meaning	Value
$MAX_STRING_LITERAL A string literal of size MAX_IN_LEN, including the quote characters.	'"' & (253 * 'A') & '"'
$MIN_INT A universal integer literal whose value is SYSTEM.MIN_INT.	-2147483648
$MIN_TASK_SIZE An integer literal whose value is the number of bits required to hold a task object which has no entries, no declarations, and NULL;" as the only statement in its body.	32
$NAME A name of a predefined numeric type other than FLOAT, INTEGER, SHORT_FLOAT, SHORT_INTEGER, LONG_FLOAT, or LONG_INTEGER.	NO_SUCH_TYPE
$NAME_LIST A list of enumeration literals in the type SYSTEM.NAME, separated by commas.	I_80X86
$NEG_BASED_INT A based integer literal whose highest order nonzero bit falls in the sign bit position of the representation for SYSTEM.MAX_INT.	16#FFFFFFFE#
$NEW_MEM_SIZE An integer literal whose value is a permitted argument for pragma memory_size, other than DEFAULT_MEM_SIZE. If there is no other value, then use DEFAULT_MEM_SIZE.	655360
$NEW_STOR_UNIT An integer literal whose value is a permitted argument for pragma storage_unit, other than DEFAULT_STOR_UNIT. If there is no other permitted value, then use value of SYSTEM.STORAGE_UNIT.	8

Name and Meaning	Value
$NEW_SYS_NAME A value of the type SYSTEM.NAME, other than $DEFAULT_SYS_NAME. If there is only one value of that type, then use that value.	I_80X86
$TASK_SIZE An integer literal whose value is the number of bits required to hold a task object which has a single entry with one inout parameter.	32
$TICK A real literal whose value is SYSTEM.TICK.	1.0/10.0

APPENDIX C

WITHDRAWN TESTS

Some tests are withdrawn from the ACVC because they do not conform to the Ada Standard. The following 44 tests had been withdrawn at the time of validation testing for the reasons indicated. A reference of the form AI-ddddd is to an Ada Commentary.

E28005C
This test expects that the string "-- TOP OF PAGE. --63" of line 204 will appear at the top of the listing page due to a pragma PAGE in line 203; but line 203 contains text that follows the pragma, and it is this that must appear at the top of the page.

A39005G
This test unreasonably expects a component clause to pack an array component into a minimum size (line 30).

B97102E
This test contains an unitended illegality: a select statement contains a null statement at the place of a selective wait alternative (line 31).

C97116A
This test contains race conditions, and it assumes that guards are evaluated indivisibly. A conforming implementation may use interleaved execution in such a way that the evaluation of the guards at lines 50 & 54 and the execution of task CHANGING_OF_THE_GUARD results in a call to REPORT.FAILED at one of lines 52 or 56.

BC3009B
This test wrongly expects that circular instantiations will be detected in several compilation units even though none of the units is illegal with respect to the units it depends on; by AI-00256, the illegality need not be detected until execution is attempted (line 95).

CD2A62D
This test wrongly requires that an array object's size be no greater than 10 although its subtype's size was specified to be 40 (line 137).

CD2A63A..D, CD2A66A..D, CD2A73A..D, CD2A76A..D [16 tests]
These tests wrongly attempt to check the size of objects of a derived type (for which a 'SIZE length clause is given) by passing them to a derived subprogram (which implicitly converts them to the parent type (Ada standard 3.4:14)). Additionally, they use the 'SIZE length clause and attribute, whose interpretation is considered problematic by the WG9 ARG.

CD2A81G, CD2A83G, CD2A84N & M, & CD50110 [5 tests]
These tests assume that dependent tasks will terminate while the main program executes a loop that simply tests for task termination; this is not the case, and the main program may loop indefinitely (lines 74, 85, 86 & 96, 86 & 96, and 58, resp.).

CD2B15C & CD7205C
These tests expect that a 'STORAGE_SIZE length clause provides precise control over the number of designated objects in a collection; the Ada standard 13.2:15 allows that such control must not be expected.

CD2D11B
This test gives a SMALL representation clause for a derived fixed-point type (at line 30) that defines a set of model numbers that are not necessarily represented in the parent type; by Commentary AI-00099, all model numbers of a derived fixed-point type must be representable values of the parent type.

CD5007B
This test wrongly expects an implicitly declared subprogram to be at the the address that is specified for an unrelated subprogram (line 303).

ED7004B, ED7005C & D, ED7006C & D [5 tests]
These tests check various aspects of the use of the three SYSTEM pragmas; the AVO withdraws these tests as being inappropriate for validation.

CD7105A
This test requires that successive calls to CALENDAR.CLOCK change by at least SYSTEM.TICK; however, by Commentary AI-00201, it is only the expected frequency of change that must be at least SYSTEM.TICK--particular instances of change may be less (line 29).

CD7203B, & CD7204B
These tests use the 'SIZE length clause and attribute, whose interpretation is considered problematic by the WG9 ARG.

CD7205D
This test checks an invalid test objective: it treats the specification of storage to be reserved for a task's activation as though it were like the specification of storage for a collection.

CE2107I
This test requires that objects of two similar scalar types be distinguished when read from a file--DATA_ERROR is expected to be raised by an attempt to read one object as of the other type. However, it is not clear exactly how the Ada standard 14.2.4:4 is to be interpreted; thus, this test objective is not considered valid. (line 90)

CE3111C
This test requires certain behavior, when two files are associated with the same external file, that is not required by the Ada standard.

CE3301A
This test contains several calls to END_OF_LINE & END_OF_PAGE that have no parameter: these calls were intended to specify a file, not to refer to STANDARD_INPUT (lines 103, 107, 118, 132, & 136).

CE3411B
This test requires that a text file's column number be set to COUNT'LAST in order to check that LAYOUT_ERROR is raised by a subsequent PUT operation. But the former operation will generally raise an exception due to a lack of available disk space, and [illegible] would thus encumber validation testing.

APPENDIX D

APPENDIX F OF THE Ada STANDARD

The only allowed implementation dependencies correspond to implementation-dependent pragmas, to certain machine-dependent conventions as mentioned in chapter 13 of the Ada Standard, and to certain allowed restrictions on representation clauses. The implementation-dependent characteristics of the AlsyCOMP_040, Version 4.3 compiler, as described in this Appendix, are provided by Alsys. Unless specifically noted otherwise, references in this appendix are to compiler documentation and not to this report. Implementation-specific portions of the package STANDARD, which are not a part of Appendix F, are:

```
package STANDARD is

 ...

 type SHORT_INTEGER is range -128 .. 127;

 type INTEGER is range -32_768 .. 32_767;

 type LONG_INTEGER is range -2_147_483_648 .. 2_147_483_647;

 type FLOAT is digits 6 range
  -2#1.111_1111_1111_1111_1111_1111#E+127
..
   2#1.111_1111_1111_1111_1111_1111#E+127;

 type LONG_FLOAT is digits 15 range
-2#1.1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111#E1023
 ..
 2#1.1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111#E1023;

 type DURATION is delta 2#0.000_000_000_000_01# range
                                            -131072.00000 .. 131071.99994;

 ...

 end STANDARD;
```

Alsys DOS to CTOS/BTOS Ada Compilation System

APPENDIX F

Version 4

Alsys Inc.
67 South Bedford Street
Burlington, MA 01803-5152

Alsys S.A.
29, Avenue de Versailles
78170 La Celle St. Cloud, France

Alsys Ltd.
Partridge House, Newtown Road
Henley-on-Thames,
Oxfordshire RG9 1EN, England

TABLE OF CONTENTS

APPENDIX F

Implementation - Dependent Characteristics

This appendix summarizes the implementation-dependent characteristics of the Alsys DOS to CTOS/BTOS Ada Compiler.

Appendix F is a required part of the *Reference Manual for the Ada Programming Language* (called the RM in this appendix).

The sections of this appendix are as follows:

1. The form, allowed places, and effect of every implementation-dependent pragma.
2. The name and the type of every implementation-dependent attribute.
3. The specification of the package SYSTEM.
4. The description of the representation clauses.
5. The conventions used for any implementation-generated name denoting implementation-dependent components.
6. The interpretation of expressions that appear in address clauses, including those for interrupts.
7. Any restrictions on unchecked conversions.
8. Any implementation-dependent characteristics of the input-output packages.
9. Characteristics of numeric types.
10. Other implementation-dependent characteristics.
11. Compiler limitations.

The name *Alsys Runtime Executive Programs* or simply *Runtime Executive* refers to the runtime library routines provided for all Ada programs. These routines implement the Ada heap, exceptions, tasking control, and other utility functions.

General systems programming notes are given in another document, the *Application Developer's Guide* (for example, parameter passing conventions needed for interface with assembly routines).

Section 1

Implementation-Dependent Pragmas

1.1 INLINE

Pragma INLINE is fully supported; however, it is not possible to inline a subprogram in a declarative part.

1.2 INTERFACE

Ada programs can interface with subprograms written in Assembler and other languages through the use of the predefined pragma INTERFACE and the implementation-defined pragma INTERFACE_NAME.

Pragma INTERFACE specifies the name of an interfaced subprogram and the name of the programming language for which parameter passing conventions will be generated. Pragma INTERFACE takes the form specified in the RM:

pragma INTERFACE (*language_name*, *subprogram_name*);

where,

- *language_name* is ASSEMBLER or ADA.
- *subprogram_name* is the name used within the Ada program to refer to the interfaced subprogram.

The only language names accepted by pragma INTERFACE are ASSEMBLER and ADA. The full implementation requirements for writing pragma INTERFACE subprograms are described in the *Application Developer's Guide.*

The language name used in the pragma INTERFACE does not have to have any relationship to the language actually used to write the interfaced subprogram. It is used only to tell the Compiler how to generate subprogram calls; that is, what kind of parameter passing techniques to use. The programmer can interface Ada programs with subroutines written in any other (compiled) language by understanding the mechanisms used for parameter passing by the Compiler and the corresponding mechanisms of the chosen external language.

1.3 INTERFACE_NAME

Pragma INTERFACE_NAME associates the name of the interfaced subprogram with the external name of the interfaced subprogram. If pragma INTERFACE_NAME is not used, then the two names are assumed to be identical. This pragma takes the form:

pragma INTERFACE_NAME (*subprogram_name*, *string_literal*);

where,

- *subprogram_name* is the name used within the Ada program to refer to the interfaced subprogram.
- *string_literal* is the name by which the interfaced subprogram is referred to at link time.

The pragma INTERFACE_NAME is used to identify routines in other languages that are not named with legal Ada identifiers. Ada identifiers can only contain letters, digits, or underscores, whereas the DOS Linker allows external names to contain other characters, for example, the dollar sign ($) or commercial at sign (@). These characters can be specified in the *string_literal* argument of the pragma INTERFACE_NAME.

The pragma INTERFACE_NAME is allowed at the same places of an Ada program as the pragma INTERFACE. (Location restrictions can be found in section 13.9 of the RM.) However, the pragma INTERFACE_NAME must always occur after the pragma INTERFACE declaration for the interfaced subprogram.

The *string_literal* of the pragma INTERFACE_NAME is passed through unchanged to the DOS object file. The maximum length of the *string_literal* is 40 characters. This limit is not checked by the Compiler, but the string is truncated by the Binder to meet the Intel object module format standard.

The user must be aware however, that some tools from other vendors do not fully support the standard object file format and may restrict the length of symbols. For example, the IBM and Microsoft assemblers silently truncate symbols at 31 characters.

The *Runtime Executive* contains several external identifiers. All such identifiers begin with either the string "ADA@" or the string "ADAS@". Accordingly, names prefixed by "ADA@" or "ADAS@" should be avoided by the user.

Example

```
package SAMPLE_DATA is
    function SAMPLE_DEVICE (X: INTEGER) return INTEGER;
    function PROCESS_SAMPLE (X: INTEGER) return INTEGER;
private
    pragma INTERFACE (ASSEMBLER, SAMPLE_DEVICE);
    pragma INTERFACE (ADA, PROCESS_SAMPLE);
    pragma INTERFACE_NAME (SAMPLE_DEVICE, "DEVIO$GET_SAMPLE");
end SAMPLE_DATA;
```

1.4 INDENT

Pragma INDENT is only used with *AdaReformat*. *AdaReformat* is the Alsys reformatter which offers the functionalities of a pretty-printer in an Ada environment.

The pragma is placed in the source file and interpreted by the Reformatter. The line

```
pragma INDENT(OFF);
```

causes *AdaReformat* not to modify the source lines after this pragma, while

```
pragma INDENT(ON);
```

causes *AdaReformat* to resume its action after this pragma.

1.5 Other Pragmas

Pragmas IMPROVE and PACK are discussed in detail in the section on representation clauses and records (Chapter 4).

Pragma PRIORITY is accepted with the range of priorities running from 1 to 10 (see the definition of the predefined package SYSTEM in Section 3). Undefined priority (no pragma PRIORITY) is treated as though it were less than any defined priority value.

In addition to pragma SUPPRESS, it is possible to suppress all checks in a given compilation by the use of the Compiler option CHECKS. (See Chapter 4 of the *User's Guide*.)

Section 2

Implementation-Dependent Attributes

2.1 P'IS_ARRAY

For a prefix P that denotes any type or subtype, this attribute yields the value TRUE if P is an array type or an array subtype; otherwise, it yields the value FALSE.

2.2 P'RECORD_DESCRIPTOR, P'ARRAY_DESCRIPTOR

These attributes are used to control the representation of implicit components of a record. (See Section 4.8 for more details.)

2.3 E'EXCEPTION_CODE

For a prefix E that denotes an exception name, this attribute yields a value that represents the internal code of the exception. The value of this attribute is of the type INTEGER.

Section 3

Specification of the package SYSTEM

The implementation does not allow the recompilation of package SYSTEM.

```
package SYSTEM is

   --      ******************************
   --      * (1) Required Definitions. *
   --      ******************************

   type NAME is (I_80x86);
   SYSTEM_NAME : constant NAME := I_80x86;

   STORAGE_UNIT : constant := 8;
   MEMORY_SIZE  : constant := 640 * 1024;

   -- System-Dependent Named Numbers:

   MIN_INT      : constant := -(2 **31);
   MAX_INT      : constant := 2**31 - 1;
   MAX_DIGITS   : constant := 15;
   MAX_MANTISSA : constant := 31;
   FINE_DELTA   : constant := 2#1.0#E-31;

   -- For the high-resolution timer, the clock resolution is
   -- 1.0 / 1024.0.
   TICK         : constant := 1.0 / 18.2;
```

```
-- Other System-Dependent Declarations:

subtype PRIORITY is INTEGER range 1 .. 10;

-- The type ADDRESS is, in fact, implemented as a
-- segment:offset pair.

type ADDRESS is private;
NULL_ADDRESS: constant ADDRESS;


--     *********************************
--     * (2) MACHINE TYPE CONVERSIONS *
--     *********************************

-- If the word / double-word operations below are used on
--    ADDRESS, then MSW yields the segment and LSW yields the
-- offset.

-- In the operations below, a BYTE_TYPE is any simple type
-- implemented on 8-bits (for example, SHORT_INTEGER), a WORD_TYPE is
-- any simple type implemented on 16-bits (for example, INTEGER), and
-- a DOUBLE_WORD_TYPE is any simple type implemented on
-- 32-bits (for example, LONG_INTEGER, FLOAT, ADDRESS).

-- Byte <==> Word conversions:

-- Get the most significant byte:
generic
   type BYTE_TYPE is private;
   type WORD_TYPE is private;
function MSB (W: WORD_TYPE) return BYTE_TYPE;
```

```
-- Get the least significant byte:
generic
   type BYTE_TYPE is private;
   type WORD_TYPE is private;
function LSB (W: WORD_TYPE) return BYTE_TYPE;

-- Compose a word from two bytes:
generic
   type BYTE_TYPE is private;
   type WORD_TYPE is private;
function WORD (MSB, LSB: BYTE_TYPE) return WORD_TYPE;

-- Word <==> Double-Word conversions:

-- Get the most significant word:
generic
   type WORD_TYPE is private;
   type DOUBLE_WORD_TYPE is private;
function MSW (W: DOUBLE_WORD_TYPE) return WORD_TYPE;

-- Get the least significant word:
generic
   type WORD_TYPE is private;
   type DOUBLE_WORD_TYPE is private;
function LSW(W: DOUBLE_WORD_TYPE) return WORD_TYPE;

-- Compose a DATA double word from two words.
generic
   type WORD_TYPE is private;
   -- The following type must be a data type
   -- (for example, LONG_INTEGER):
   type DATA_DOUBLE_WORD is private;
function DOUBLE_WORD (MSW, LSW: WORD_TYPE) return DATA_DOUBLE_WORD;
```

```
   -- Compose a REFERENCE double word from two words.
   generic
      type WORD_TYPE is private;
      -- The following type must be a reference type
      -- (for example, access or ADDRESS):
      type REF_DOUBLE_WORD is private;
   function REFERENCE (SEGMENT, OFFSET: WORD_TYPE) return REF_DOUBLE_WORD;


   --      ******************************
   --      * (3) OPERATIONS ON ADDRESS *
   --      ******************************

   -- You can get an address via 'ADDRESS attribute or by
   -- instantiating the function REFERENCE, above, with
   -- appropriate types.
   -- Some addresses are used by the Compiler.  For example,
   -- the display is located at the low end of the DS segment,
   -- and addresses SS:0 through SS:128 hold the task control
   -- block and other information.  Writing into these areas
   -- will have unpredictable results.-- Note that no operations are defined to get the
   -- values of the segment registers, but if it is necessary an
   -- interfaced function can be written.

   generic
      type OBJECT is private;
   function FETCH_FROM_ADDRESS (FROM: ADDRESS) return OBJECT;
   generic
      type OBJECT is private;
   procedure ASSIGN_TO_ADDRESS (OBJ: OBJECT; TO: ADDRESS);

private

   ...

end SYSTEM;
```

Section 4

Support for Representation Clauses

This section explains how objects are represented and allocated by the Alsys CTOS/BTOS Ada compiler and how it is possible to control this using representation clauses. Applicable restrictions on representation clauses are also described.

The representation of an object is closely connected with its type. For this reason this section addresses successively the representation of enumeration, integer, floating point, fixed point, access, task, array and record types. For each class of type the representation of the corresponding objects is described.

Except in the case of array and record types, the description for each class of type is independent of the others. To understand the representation of array and record types it is necessary to understand first the representation of their components.

Apart from implementation defined pragmas, Ada provides three means to control the size of objects:

- a (predefined) pragma PACK, applicable to array types
- a record representation clause
- a size specification

For each class of types the effect of a size specification is described. Interactions among size specifications, packing and record representation clauses is described under the discussion of array and record types.

Representation clauses on derived record types or derived tasks types are not supported.

Size representation clauses on types derived from private types are not supported when the derived type is declared outside the private part of the defining package.

4.1 Enumeration Types

4.1.1 Enumeration Literal Encoding

When no enumeration representation clause applies to an enumeration type, the internal code associated with an enumeration literal is the position number of the enumeration literal. Then, for an enumeration type with n elements, the internal codes are the integers 0, 1, 2, .. , n-1.

An enumeration representation clause can be provided to specify the value of each internal code as described in RM 13.3. The Alsys compiler fully implements enumeration representation clauses.

As internal codes must be machine integers the internal codes provided by an enumeration representation clause must be in the range $-2^{31} .. 2^{31}-1$.

An enumeration value is always represented by its internal code in the program generated by the compiler.

4.1.2 Enumeration Types and Object Sizes

Minimum size of an enumeration subtype

The minimum possible size of an enumeration subtype is the minimum number of bits that is necessary for representing the internal codes of the subtype values in normal binary form.

A static subtype, with a null range has a minimum size of 1. Otherwise, if m and M are the values of the internal codes associated with the first and last enumeration values of the subtype, then its minimum size L is determined as follows. For $m >= 0$, L is the smallest positive integer such that $M <= 2^{L}-1$. For $m < 0$, L is the smallest positive integer such that $-2^{L-1} <= m$ and $M <= 2^{L-1}-1$. For example:

```
type COLOR is (GREEN, BLACK, WHITE, RED, BLUE, YELLOW);
-- The minimum size of COLOR is 3 bits.

subtype BLACK_AND_WHITE is COLOR range BLACK .. WHITE;
-- The minimum size of BLACK_AND_WHITE is 2 bits.
```

```
subtype BLACK_OR_WHITE is BLACK_AND_WHITE range X .. X;
-- Assuming that X is not static, the minimum size of BLACK_OR_WHITE is
-- 2 bits (the same as the minimum size of its type mark BLACK_AND_WHITE).
```

Size of an enumeration subtype

When no size specification is applied to an enumeration type or first named subtype, the objects of that type or first named subtype are represented as signed machine integers. The machine provides 8, 16 and 32 bit integers, and the compiler selects automatically the smallest signed machine integer which can hold each of the internal codes of the enumeration type (or subtype). The size of the enumeration type and of any of its subtypes is thus 8, 16 or 32 bits.

When a size specification is applied to an enumeration type, this enumeration type and each of its subtypes has the size specified by the length clause. The same rule applies to a first named subtype. The size specification must of course specify a value greater than or equal to the minimum size of the type or subtype to which it applies:

```
type EXTENDED is
    (-- The usual ASCII character set.
    NUL, SOH, STX, ETX, EOT, ENQ, ACK, BEL,
    ...
    'x', 'y', 'z', '{', '|', '}', '~', DEL,

     -- Extended characters
    C_CEDILLA_CAP, U_UMLAUT, E_ACUTE, ...);

for EXTENDED'SIZE use 8;
-- The size of type EXTENDED will be one byte. Its objects will be represented
-- as unsigned 8 bit integers.
```

The Alsys compiler fully implements size specifications. Nevertheless, as enumeration values are coded using integers, the specified length cannot be greater than 32 bits.

Size of the objects of an enumeration subtype

Provided its size is not constrained by a record component clause or a pragma PACK, an object of an enumeration subtype has the same size as its subtype.

4.2 Integer Types

There are three predefined integer types in the Alsys implementation for I80x86 machines:

```
type SHORT_INTEGER        is range -2**07 .. 2**07-1;
type INTEGER              is range -2**15 .. 2**15-1;
type LONG_INTEGER         is range -2**31 .. 2**31-1;
```

4.2.1 Integer Type Representation

An integer type declared by a declaration of the form:

```
type T is range L .. R;
```

is implicitly derived from a predefined integer type. The compiler automatically selects the predefined integer type whose range is the smallest that contains the values L to R inclusive.

Binary code is used to represent integer values. Negative numbers are represented using two's complement.

4.2.2 Integer Type and Object Size

Minimum size of an integer subtype

The minimum possible size of an integer subtype is the minimum number of bits that is necessary for representing the internal codes of the subtype values in normal binary form.

For a static subtype, if it has a null range its minimum size is 1. Otherwise, if m and M are the lower and upper bounds of the subtype, then its minimum size L is determined as follows. For $m >= 0$, L is the smallest positive integer such that $M <= 2^{L-1}$. For $m < 0$, L is the smallest positive integer that $-2^{L-1} <= m$ and $M <= 2^{L-1}-1$. For example:

```
subtype S is INTEGER range 0 .. 7;
-- The minimum size of S is 3 bits.
```

```
subtype D is S range X .. Y;
-- Assuming that X and Y are not static, the minimum size of
-- D is 3 bits (the same as the minimum size of its type mark S).
```

Size of an integer subtype

The sizes of the predefined integer types SHORT_INTEGER, INTEGER and LONG_INTEGER are respectively 8, 16 and 32 bits.

When no size specification is applied to an integer type or to its first named subtype (if any), its size and the size of any of its subtypes is the size of the predefined type from which it derives, directly or indirectly. For example:

```
type S is range 80 .. 100;
-- S is derived from SHORT_INTEGER, its size is 8 bits.

type J is range 0 .. 255;
-- J is derived from INTEGER, its size is 16 bits.

type N is new J range 80 .. 100;
-- N is indirectly derived from INTEGER, its size is 16 bits.
```

When a size specification is applied to an integer type, this integer type and each of its subtypes has the size specified by the length clause. The same rule applies to a first named subtype. The size specification must of course specify a value greater than or equal to the minimum size of the type or subtype to which it applies:

```
type S is range 80 .. 100;
for S'SIZE use 32;
-- S is derived from SHORT_INTEGER, but its size is 32 bits
-- because of the size specification.

type J is range 0 .. 255;
for J'SIZE use 8;
-- J is derived from INTEGER, but its size is 8 bits because
-- of the size specification.

type N is new J range 80 .. 100;
-- N is indirectly derived from INTEGER, but its size is
-- 8 bits because N inherits the size specification of J.
```

Size of the objects of an integer subtype

Provided its size is not constrained by a record component clause or a pragma PACK, an object of an integer subtype has the same size as its subtype.

4.3 Floating Point Types

There are two predefined floating point types in the Alsys implementation for I80x86 machines:

```
type FLOAT is
    digits 6 range -(2.0 - 2.0**(-23))*2.0**127 .. (2.0 - 2.0**(-23))*2.0**127;

type LONG_FLOAT is
    digits 15 range -(2.0 - 2.0**(-51))*2.0**1023 .. (2.0 - 2.0**(-51))*2.0**1023;
```

4.3.1 Floating Point Type Representation

A floating point type declared by a declaration of the form:

```
type T is digits D [range L .. R];
```

is implicitly derived from a predefined floating point type. The compiler automatically selects the smallest predefined floating point type whose number of digits is greater than or equal to D and which contains the values L to R inclusive.

In the program generated by the compiler, floating point values are represented using the IEEE standard formats for single and double floats.

The values of the predefined type FLOAT are represented using the single float format. The values of the predefined type LONG_FLOAT are represented using the double float format. The values of any other floating point type are represented in the same way as the values of the predefined type from which it derives, directly or indirectly.

4.3.2 Floating Point Type and Object Size

The minimum possible size of a floating point subtype is 32 bits if its base type is FLOAT or a type derived from FLOAT; it is 64 bits if its base type is LONG_FLOAT or a type derived from LONG_FLOAT.

The sizes of the predefined floating point types FLOAT and LONG_FLOAT are respectively 32 and 64 bits.

The size of a floating point type and the size of any of its subtypes is the size of the predefined type from which it derives directly or indirectly.

The only size that can be specified for a floating point type or first named subtype using a size specification is its usual size (32 or 64 bits).

An object of a floating point subtype has the same size as its subtype.

4.4 Fixed Point Types

4.4.1 Fixed Point Type Representation

If no specification of small applies to a fixed point type, then the value of small is determined by the value of delta as defined by RM 3.5.9.

A specification of small can be used to impose a value of small. The value of small is required to be a power of two.

To implement fixed point types, the Alsys compiler for I80x86 machines uses a set of anonymous predefined types of the form:

```
type SHORT_FIXED is delta D range (-2.0**7-1)*S .. 2.0**7*S;
for SHORT_FIXED'SMALL use S;

type FIXED is delta D range (-2.0**15-1)*S .. 2.0**15*S;
for FIXED'SMALL use S;

type LONG_FIXED is delta D range (-2.0**31-1)*S .. 2.0**31*S;
for LONG_FIXED'SMALL use S;
```

where D is any real value and S any power of two less than or equal to D.

A fixed point type declared by a declaration of the form:

```
type T is delta D range L .. R;
```

possibly with a small specification:

```
for T'SMALL use S;
```

is implicitly derived from a predefined fixed point type. The compiler automatically selects the predefined fixed point type whose small and delta are the same as the small and delta of T and whose range is the shortest that includes the values L to R inclusive.

In the program generated by the compiler, a safe value V of a fixed point subtype F is represented as the integer:

```
V / F'BASE'SMALL
```

4.4.2 Fixed Point Type and Object Size

Minimum size of a fixed point subtype

The minimum possible size of a fixed point subtype is the minimum number of binary digits that is necessary for representing the values of the range of the subtype using the small of the base type.

For a static subtype, if it has a null range its minimum size is 1. Otherwise, s and S being the bounds of the subtype, if i and I are the integer representations of m and M, the smallest and the greatest model numbers of the base type such that s < m and M < S, then the minimum size L is determined as follows. For i >= 0, L is the smallest positive integer such that $I <= 2^{L-1}$. For i < 0, L is the smallest positive integer such that $-2^{L-1} <= i$ and $I <= 2^{L-1}-1$.

```
type F is delta 2.0 range 0.0 .. 500.0;
    -- The minimum size of F is 8 bits.

subtype S is F delta 16.0 range 0.0 .. 250.0;
    -- The minimum size of S is 7 bits.

subtype D is S range X .. Y;
-- Assuming that X and Y are not static, the minimum size of D is 7 bits
-- (the same as the minimum size of its type mark S).
```

Size of a fixed point subtype

The sizes of the predefined fixed point types SHORT_FIXED. FIXED and LONG_FIXED are respectively 8, 16 and 32 bits.

When no size specification is applied to a fixed point type or to its first named subtype, its size and the size of any of its subtypes is the size of the predefined type from which it derives directly or indirectly. For example:

```
type S is delta 0.01 range 0.8 .. 1.0;
-- S is derived from an 8 bit predefined fixed type, its size is 8 bits.

type F is delta 0.01 range 0.0 .. 2.0;
-- F is derived from a 16 bit predefined fixed type, its size is 16 bits.

type N is new F range 0.8 .. 1.0;
-- N is indirectly derived from a 16 bit predefined fixed type, its size is 16 bits.
```

When a size specification is applied to a fixed point type, this fixed point type and each of its subtypes has the size specified by the length clause. The same rule applies to a first named subtype. The size specification must of course specify a value greater than or equal to the minimum size of the type or subtype to which it applies:

```
type S is delta 0.01 range 0.8 .. 1.0;
for S'SIZE use 32;
-- S is derived from an 8 bit predefined fixed type, but its size is 32 bits
-- because of the size specification.

type F is delta 0.01 range 0.0 .. 2.0;
for F'SIZE use 8;
-- F is derived from a 16 bit predefined fixed type, but its size is 8 bits
-- because of the size specification.

type N is new F range 0.8 .. 1.0;
-- N is indirectly derived from a 16 bit predefined fixed type, but its size is
-- 8 bits because N inherits the size specification of F.
```

The Alsys compiler fully implements size specifications. Nevertheless, as fixed point objects are represented using machine integers, the specified length cannot be greater than 32 bits.

Size of the objects of a fixed point subtype

Provided its size is not constrained by a record component clause or a pragma PACK, an object of a fixed point type has the same size as its subtype.

4.5 Access Types and Collections

Access Types and Objects of Access Types

The only size that can be specified for an access type using a size specification is its usual size (32 bits).

An object of an access subtype has the same size as its subtype, thus an object of an access subtype is always 32 bits long.

Collection Size

As described in RM 13.2, a specification of collection size can be provided in order to reserve storage space for the collection of an access type.

When no STORAGE_SIZE specification applies to an access type, no storage space is reserved for its collection, and the value of the attribute STORAGE_SIZE is then 0. The maximum size allowed for a collection is 64k bytes.

4.6 Task Types

Storage for a task activation

As described in RM 13.2, a length clause can be used to specify the storage space (that is, the stack size) for the activation of each of the tasks of a given type. Alsys also allows the task stack size, for all tasks, to be established using a Binder option. If a length clause is given for a task type, the value indicated at bind time is ignored for this task type, and the length clause is obeyed. When no length clause is used to specify the storage space to be reserved for a task activation, the storage space indicated at bind time is used for this activation.

A length clause may not be applied to a derived task type. The same storage space is reserved for the activation of a task of a derived type as for the activation of a task of the parent type.

The minimum size of a task subtype is 32 bits.

A size specification has no effect on a task type. The only size that can be specified using such a length clause is its usual size (32 bits).

An object of a task subtype has the same size as its subtype. Thus an object of a task subtype is always 32 bits long.

4.7 Array Types

Each array is allocated in a contiguous area of storage units. All the components have the same size. A gap may exist between two consecutive components (and after the last one). All the gaps have the same size.

4.7.1 Array Layout and Structure and Pragma PACK

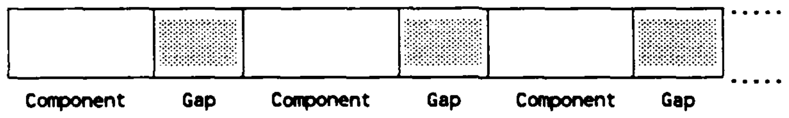

If pragma PACK is not specified for an array, the size of the components is the size of the subtype of the components:

```
type A is array (1 .. 8) of BOOLEAN;
-- The size of the components of A is the size of the type BOOLEAN: 8 bits.

type DECIMAL_DIGIT is range 0 .. 9;
for DECIMAL_DIGIT'SIZE use 4;
type BINARY_CODED_DECIMAL is
    array (INTEGER range < >) of DECIMAL_DIGIT;
-- The size of the type DECIMAL_DIGIT is 4 bits. Thus in an array of
-- type BINARY_CODED_DECIMAL each component will be represented on
-- 4 bits as in the usual BCD representation.
```

If pragma PACK is specified for an array and its components are neither records nor arrays, the size of the components is the minimum size of the subtype of the components:

```
type A is array (1 .. 8) of BOOLEAN;
pragma PACK(A);
-- The size of the components of A is the minimum size of the type BOOLEAN:
-- 1 bit.

type DECIMAL_DIGIT is range 0 .. 9;
for DECIMAL_DIGIT'SIZE use 32;
type BINARY_CODED_DECIMAL is
    array (INTEGER range < >) of DECIMAL_DIGIT;
pragma PACK(BINARY_CODED_DECIMAL);
-- The size of the type DECIMAL_DIGIT is 32 bits, but, as
-- BINARY_CODED_DECIMAL is packed, each component of an array of this
-- type will be represented on 4 bits as in the usual BCD representation.
```

Packing the array has no effect on the size of the components when the components are records or arrays, since records and arrays may be assigned addresses consistent with the alignment of their subtypes.

Gaps

If the components are records or arrays, no size specification applies to the subtype of the components and the array is not packed, then the compiler may choose a representation with a gap after each component; the aim of the insertion of such gaps is to optimize access to the array components and to their subcomponents. The size of the gap is chosen so that the relative displacement of consecutive components is a multiple of the alignment of the subtype of the components. This strategy allows each component and subcomponent to have an address consistent with the alignment of its subtype:

```
type R is
    record
        K : INTEGER;
        B : BOOLEAN;
    end record;
for R use
    record
        K at 0 range 0 .. 15;
        B at 2 range 0 .. 0;
    end record;
-- Record type R is byte aligned. Its size is 17 bits.
```

```
type A is array (1 .. 10) of R;
-- A gap of 7 bits is inserted after each component in order to respect the
-- alignment of type R. The size of an array of type A will be 240 bits.
```

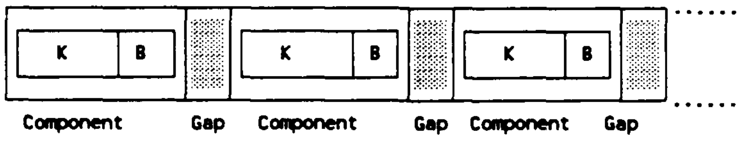

Array of type A: each subcomponent K has an even offset.

If a size specification applies to the subtype of the components or if the array is packed, no gaps are inserted:

```
type R is
    record
        K : INTEGER;
        B : BOOLEAN;
    end record;

type A is array (1 .. 10) of R;
pragma PACK(A);
-- There is no gap in an array of type A because A is packed.
-- The size of an object of type A will be 270 bits.

type NR is new R;
for NR'SIZE use 24;

type B is array (1 .. 10) of NR;
-- There is no gap in an array of type B because
-- NR has a size specification.
-- The size of an object of type B will be 240 bits.
```

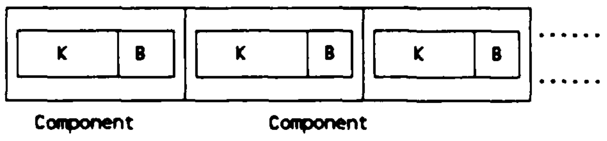

Array of type A or B

4.7.2 Array Subtype and Object Size

Size of an array subtype

The size of an array subtype is obtained by multiplying the number of its components by the sum of the size of the components and the size of the gaps (if any). If the subtype is unconstrained, the maximum number of components is considered.

The size of an array subtype cannot be computed at compile time

- if it has non-static constraints or is an unconstrained array type with non-static index subtypes (because the number of components can then only be determined at run time).
- if the components are records or arrays and their constraints or the constraints of their subcomponents (if any) are not static (because the size of the components and the size of the gaps can then only be determined at run time).

As has been indicated above, the effect of a pragma PACK on an array type is to suppress the gaps. The consequence of packing an array type is thus to reduce its size.

If the components of an array are records or arrays and their constraints or the constraints of their subcomponents (if any) are not static, the compiler ignores any pragma PACK applied to the array type but issues a warning message. Apart from this limitation, array packing is fully implemented by the Alsys compiler.

A size specification applied to an array type or first named subtype has no effect. The only size that can be specified using such a length clause is its usual size. Nevertheless, such a length clause can be useful to verify that the layout of an array is as expected by the application.

Size of the objects of an array subtype

The size of an object of an array subtype is always equal to the size of the subtype of the object.

4.8 Record Types

4.8.1 Basic Record Structure

Layout of a record

Each record is allocated in a contiguous area of storage units. The size of a record component depends on its type.

The positions and the sizes of the components of a record type object can be controlled using a record representation clause as described in RM 13.4. In the Alsys implementation for I80x86 machines there is no restriction on the position that can be specified for a component of a record. If a component is not a record or an array, its size can be any size from the minimum size to the size of its subtype. If a component is a record or an array, its size must be the size of its subtype.

```
type DEVICE_INFO_RECORD is
   record
      BIT15   : BOOLEAN;   -- Bit 15 (reserved)
      CTRL    : BOOLEAN;   -- Bit 14 (true if control strings processed)
      NETWORK : BOOLEAN;   -- Bit 13 (true if device is on network)
      BIT12   : BOOLEAN;   -- Bit 12 (reserved)
      BIT11   : BOOLEAN;   -- Bit 11 (reserved)
      BIT10   : BOOLEAN;   -- Bit 10 (reserved)
      BIT9    : BOOLEAN;   -- Bit 9 (reserved)
      BIT8    : BOOLEAN;   -- Bit 8 (reserved)
      ISDEV   : BOOLEAN;   -- Bit 7 (true if device, false if disk file)
      EOF     : BOOLEAN;   -- Bit 6 (true if at end of file)
      BINARY  : BOOLEAN;   -- Bit 5 (true if binary (raw) mode)
      BIT4    : BOOLEAN;   -- Bit 4 (reserved)
      ISCLK   : BOOLEAN;   -- Bit 3 (true if clock device)
      ISNUL   : BOOLEAN;   -- Bit 2 (true if NUL device)
      ISCOT   : BOOLEAN;   -- Bit 1 (true if console output device)
      ISCIN   : BOOLEAN;   -- Bit 0 (true if console input device)
   end record;
```

```
for DEVICE_INFO_RECORD use
   record
      BIT15    at 1 range 7 .. 7;    -- Bit 15
      CTRL     at 1 range 6 .. 6;    -- Bit 14
      NETWORK  at 1 range 5 .. 5;    -- Bit 13
      BIT12    at 1 range 4 .. 4;    -- Bit 12
      BIT11    at 1 range 3 .. 3;    -- Bit 11
      BIT10    at 1 range 2 .. 2;    -- Bit 10
      BIT9     at 1 range 1 .. 1;    -- Bit 9
      BIT8     at 1 range 0 .. 0;    -- Bit 8
      ISDEV    at 0 range 7 .. 7;    -- Bit 7
      EOF      at 0 range 6 .. 6;    -- Bit 6
      BINARY   at 0 range 5 .. 5;    -- Bit 5
      BIT4     at 0 range 4 .. 4;    -- Bit 4
      ISCLK    at 0 range 3 .. 3;    -- Bit 3
      ISNUL    at 0 range 2 .. 2;    -- Bit 2
      ISCOT    at 0 range 1 .. 1;    -- Bit 1
      ISCIN    at 0 range 0 .. 0;    -- Bit 0
   end record;
```

Pragma PACK has no effect on records. *It is unnecessary because record representation* clauses provide full control over record layout.

A record representation clause need not specify the position and the size for every component. If no component clause applies to a component of a record, its size is the size of its subtype.

4.8.2 Indirect Components

If the offset of a component cannot be computed at compile time, this offset is stored in the record objects at run time and used to access the component. Such a component is said to be indirect while other components are said to be direct:

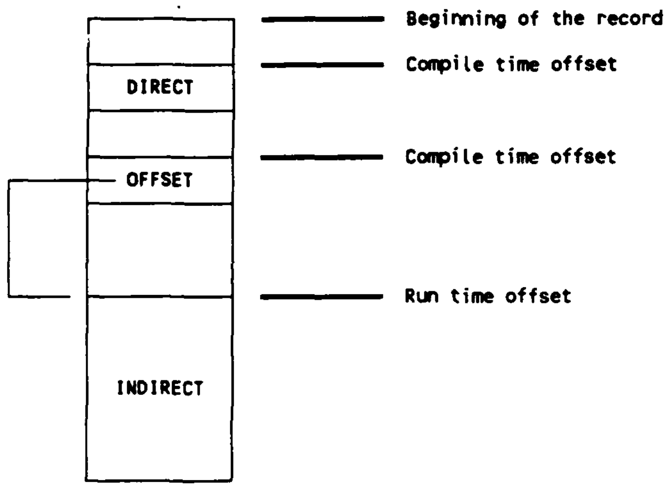

A direct and an indirect component

If a record component is a record or an array, the size of its subtype may be evaluated at run time and may even depend on the discriminants of the record. We will call these components dynamic components:

```
type DEVICE is (SCREEN, PRINTER);

type COLOR is (GREEN, RED, BLUE);

type SERIES is array (POSITIVE range < >) of INTEGER;

type GRAPH (L : NATURAL) is
    record
        X : SERIES(1 .. L); -- The size of X depends on L
        Y : SERIES(1 .. L); -- The size of Y depends on L
    end record;

Q : POSITIVE;
```

```
type PICTURE (N : NATURAL; D : DEVICE) is
   record
      F : GRAPH(N); -- The size of F depends on N
      S : GRAPH(Q); -- The size of S depends on Q
      case D is
        when SCREEN =>
          C : COLOR;
        when PRINTER =>
          null;
      end case;
   end record;
```

Any component placed after a dynamic component has an offset which cannot be evaluated at compile time and is thus indirect. In order to minimize the number of indirect components, the compiler groups the dynamic components together and places them at the end of the record:

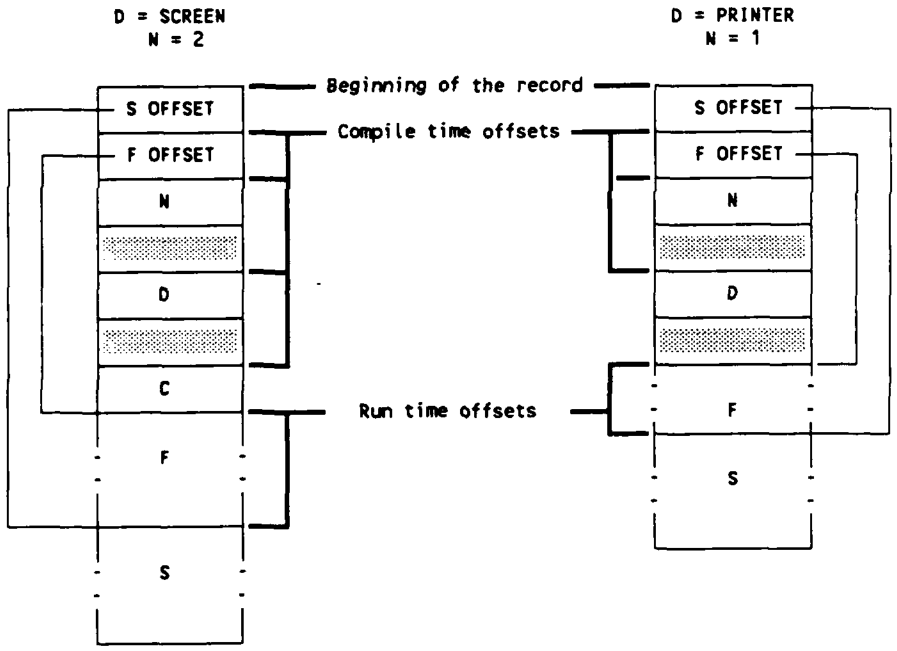

The record type PICTURE: F and S are placed at the end of the record

Note that Ada does not allow representation clauses for record components with non-static bounds [RM 13.4.7], so the compiler's grouping of dynamic components does not conflict with the use of representation clauses.

Because of this approach, the only indirect components are dynamic components. But not all dynamic components are necessarily indirect: if there are dynamic components in a component list which is not followed by a variant part, then exactly one dynamic component of this list is a direct component because its offset can be computed at compilation time (the only dynamic components that are direct components are in this situation):

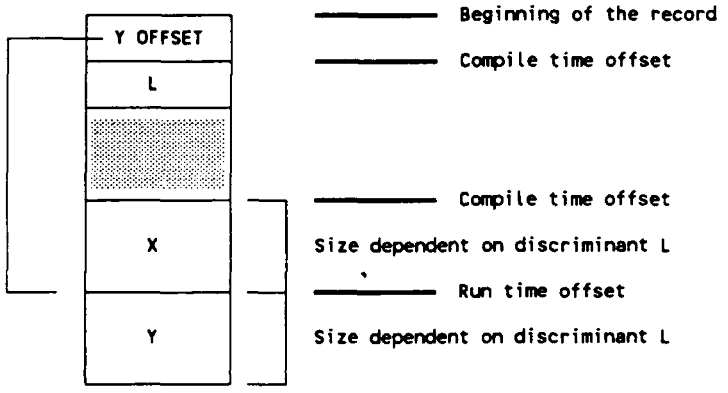

The record type GRAPH: the dynamic component X is a direct component.

The offset of an indirect component is always expressed in storage units.

The space reserved for the offset of an indirect component must be large enough to store the size of any value of the record type (the maximum potential offset). The compiler evaluates an upper bound MS of this size and treats an offset as a component having an anonymous integer type whose range is 0 .. MS.

If C is the name of an indirect component, then the offset of this component can be denoted in a component clause by the implementation generated name C'OFFSET.

4.8.3 Implicit Components

In some circumstances, access to an object of a record type or to its components involves computing information which only depends on the discriminant values. To avoid recomputation (which would degrade performance) the compiler stores this information in the record objects, updates it when the values of the discriminants are modified and uses it when the objects or its components are accessed. This information is stored in special components called implicit components.

An implicit component may contain information which is used when the record object or several of its components are accessed. In this case the component will be included in any record object (the implicit component is considered to be declared before any variant part in the record type declaration). There can be two components of this kind; one is called RECORD_SIZE and the other VARIANT_INDEX.

On the other hand an implicit component may be used to access a given record component. In that case the implicit component exists whenever the record component exists (the implicit component is considered to be declared at the same place as the record component). Components of this kind are called ARRAY_DESCRIPTORs or RECORD_DESCRIPTORs.

RECORD_SIZE

This implicit component is created by the compiler when the record type has a variant part and its discriminants are defaulted. It contains the size of the storage space necessary to store the current value of the record object (note that the storage effectively allocated for the record object may be more than this).

The value of a RECORD_SIZE component may denote a number of bits or a number of storage units. In general it denotes a number of storage units, but if any component clause specifies that a component of the record type has an offset or a size which cannot be expressed using storage units, then the value designates a number of bits.

The implicit component RECORD_SIZE must be large enough to store the maximum size of any value of the record type. The compiler evaluates an upper bound MS of this size and then considers the implicit component as having an anonymous integer type whose range is 0 .. MS.

If R is the name of the record type, this implicit component can be denoted in a component clause by the implementation generated name R'RECORD_SIZE. This allows user control over the position of the implicit component in the record.

VARIANT_INDEX

This implicit component is created by the compiler when the record type has a variant part. It indicates the set of components that are present in a record value. It is used when a discriminant check is to be done.

Component lists in variant parts that themselves do not contain a variant part are numbered. These numbers are the possible values of the implicit component VARIANT_INDEX.

```
type VEHICLE is (AIRCRAFT, ROCKET, BOAT, CAR);

type DESCRIPTION (KIND : VEHICLE := CAR) is
    record
        SPEED : INTEGER;
        case KIND is
          when AIRCRAFT | CAR =>
            WHEELS : INTEGER;
            case KIND is
              when AIRCRAFT =>           -- 1
                WINGSPAN : INTEGER;
              when others =>       -- 2
                null;
            end case;
          when BOAT => -- 3
            STEAM  : BOOLEAN;
          when ROCKET =>         -- 4
            STAGES : INTEGER;
        end case;
end record;
```

The value of the variant index indicates the set of components that are present in a record value:

Variant Index	Set
1	{KIND, SPEED, WHEELS, WINGSPAN}
2	{KIND, SPEED, WHEELS}
3	{KIND, SPEED, STEAM}
4	{KIND, SPEED, STAGES}

A comparison between the variant index of a record value and the bounds of an interval is enough to check that a given component is present in the value:

Component	Interval
KIND	--
SPEED	--
WHEELS	1 .. 2
WINGSPAN	1 .. 1
STEAM	3 .. 3
STAGES	4 .. 4

The implicit component VARIANT_INDEX must be large enough to store the number V of component lists that don't contain variant parts. The compiler treats this implicit component as having an anonymous integer type whose range is 1 .. V.

If R is the name of the record type, this implicit component can be denoted in a component clause by the implementation generated name R'VARIANT_INDEX. This allows user control over the position of the implicit component in the record.

ARRAY_DESCRIPTOR

An implicit component of this kind is associated by the compiler with each record component whose subtype is an anonymous array subtype that depends on a discriminant of the record. It contains information about the component subtype.

The structure of an implicit component of kind ARRAY_DESCRIPTOR is not described in this documentation. Nevertheless, if a programmer is interested in specifying the location of a component of this kind using a component clause, size of the component may be obtained using the ASSEMBLY parameter in the COMPILE command.

The compiler treats an implicit component of the kind ARRAY_DESCRIPTOR as having an anonymous array type. If C is the name of the record component whose subtype is described by the array descriptor, then this implicit component can be denoted in a component clause by the implementation generated name C'ARRAY_DESCRIPTOR. This allows user control over the position of the implicit component in the record.

RECORD_DESCRIPTOR

An implicit component of this kind is associated by the compiler with each record component whose subtype is an anonymous record subtype that depends on a discriminant of the record. It contains information about the component subtype.

The structure of an implicit component of kind RECORD_DESCRIPTOR is not described in this documentation. Nevertheless, if a programmer is interested in specifying the location of a component of this kind using a component clause, the size of the component may be obtained using the ASSEMBLY parameter in the COMPILE command.

The compiler treats an implicit component of the kind RECORD_DESCRIPTOR as having an anonymous array type. If C is the name of the record component whose subtype is described by the record descriptor, then this implicit component can be denoted in a component clause by the implementation generated name C'RECORD_DESCRIPTOR. This allows user control over the position of the implicit component in the record.

Suppression of Implicit Components

The Alsys implementation provides the capability of suppressing the implicit components **pragma** IMPROVE (TIME | SPACE , [ON =>] simple_name);

The first argument specifies whether TIME or SPACE is the primary criterion for the choice of the representation of the record type that is denoted by the second argument.

If TIME is specified, the compiler inserts implicit components as described above. If on the other hand SPACE is specified, the compiler only inserts a VARIANT_INDEX or a RECORD_SIZE component if this component appears in a record representation clause that applies to the record type. A record representation clause can thus be used to keep one implicit component while suppressing the other.

A pragma IMPROVE that applies to a given record type can occur anywhere that a representation clause is allowed for this type.

4.8.4 Size of Record Types and Objects

Size of a record subtype

Unless a component clause specifies that a component of a record type has an offset or a size which cannot be expressed using storage units, the size of a record subtype is rounded up to a whole number of storage units.

The size of a constrained record subtype is obtained by adding the sizes of its components and the sizes of its gaps (if any). This size is not computed at compile time

- when the record subtype has non-static constraints,
- when a component is an array or a record and its size is not computed at compile time.

The size of an unconstrained record subtype is obtained by adding the sizes of the components and the sizes of the gaps (if any) of its largest variant. If the size of a component or of a gap cannot be evaluated exactly at compile time an upper bound of this size is used by the compiler to compute the subtype size.

A size specification applied to a record type or first named subtype has no effect. The only size that can be specified using such a length clause is its usual size. Nevertheless, such a length clause can be useful to verify that the layout of a record is as expected by the application.

Size of an object of a record subtype

An object of a constrained record subtype has the same size as its subtype.

An object of an unconstrained record subtype has the same size as its subtype if this size is less than or equal to 8 kb. If the size of the subtype is greater than this, the object has the size necessary to store its current value; storage space is allocated and released as the discriminants of the record change.

Section 5

Conventions for Implementation-Generated Names

The Compiler may add fields to record objects and have descriptors in memory for record or array objects. These fields are accessible to the user through implementation-generated attributes (See Section 2.3).

The following predefined packages are reserved to Alsys and cannot be recompiled in Version 4.2:

```
ALSYS_BASIC_IO
ALSYS_ADA_RUNTIME
ALSYS_BASIC_IO
ALSYS_BASIC_DIRECT_IO
ALSYS_BASIC_SEQUENTIAL_IO
```

Section 6

Address Clauses

6.1 Address Clauses for Objects

An address clause can be used to specify an address for an object as described in RM 13.5. When such a clause applies to an object the compiler does not cause storage to be allocated for the object. The program accesses the object using the address specified in the clause. It is the responsibility of the user therefore to make sure that a valid allocation of storage has been done at the specified address.

An address clause is not allowed for task objects, for unconstrained records whose size is greater than 8k bytes or for a constant.

There are a number of ways to compose a legal address expression for use in an address clause. The most direct ways are:

- For the case where the memory is defined in Ada as another object, use the 'ADDRESS attribute to obtain the argument for the address clause for the second object.
- For the case where an absolute address is known to the programmer, instantiate the generic function SYSTEM.REFERENCE on a 16 bit unsigned integer type (either from package UNSIGNED, or by use of a length clause on a derived integer type or subtype) and on type SYSTEM.ADDRESS. Then the values of the desired segment and offset can be passed as the actual parameters to the instantiated function in the simple expression part of the address clause. See Section 3 for the specification of package SYSTEM.
- For the case where the desired location is memory defined in assembly or another non-Ada language (is relocatable), an interfaced routine may be used to obtain the appropriate address from referencing information known to the other language.

In all cases other than the use of an address attribute, the programmer must ensure that the segment part of the argument is a selector if the program is to run in protected mode. Refer to the *Application Developers' Guide*, Section 5.5 for more information on protected mode machine oriented programming.

6.2 Address Clauses for Program Units

Address clauses for program units are not implemented in the current version of the compiler.

6.3 Address Clauses for Entries

Address clauses for entries are not implemented in the current version of the compiler.

Section 7

Unchecked Conversions

Unchecked conversions are allowed between any types provided the instantiation of UNCHECKED_CONVERSION is legal Ada. It is the programmer's responsibility to determine if the desired effect is achieved.

If the target type has a smaller size than the source type then the target is made of the least significant bits of the source.

Section 8

Input-Output Packages

The RM defines the predefined input-output packages SEQUENTIAL_IO, DIRECT_IO, and TEXT_IO, and describes how to use the facilities available within these packages. The RM also defines the package IO_EXCEPTIONS, which specifies the exceptions that can be raised by the predefined input-output packages.

In addition the RM outlines the package LOW_LEVEL_IO, which is concerned with low-level machine-dependent input-output, such as would possibly be used to write device drivers or access device registers. LOW_LEVEL_IO has not been implemented. The use of interfaced subprograms is recommended as an alternative.

8.1 Correspondence between External Files and DOS Files

Ada input-output is defined in terms of external files. Data is read from and written to external files. Each external file is implemented as a standard DOS file, including the use of STANDARD_INPUT and STANDARD_OUTPUT.

The name of an external file can be either

- the null string
- a DOS filename
- a DOS special file or device name (for example, CON and PRN)

If the name is a null string, the associated external file is a temporary file and will cease to exist when the program is terminated. The file will be placed in the current directory and its name will be chosen by DOS.

If the name is a DOS filename, the filename will be interpreted according to standard DOS conventions (that is, relative to the current directory). The exception NAME_ERROR is raised if the name part of the filename has more than 8 characters or if the extension part has more than 3 characters.

If an existing DOS file is specified to the CREATE procedure, the contents of the file will be deleted before writing to the file.

If a non-existing directory is specified in a file path name to CREATE, the directory will not be created, and the exception NAME_ERROR is raised.

8.2 Error Handling

DOS errors are translated into Ada exceptions, as defined in the RM by package IO_EXCEPTIONS. In particular, DEVICE_ERROR is raised in cases of drive not ready, unknown media, disk full or hardware errors on the disk (such as read or write fault).

8.3 The FORM Parameter

The form parameter is a string, formed from a list of attributes, with attributes separated by commas. The string is not case sensitive. The attributes specify:

- Buffering

  BUFFER_SIZE => *size_in_bytes*

- Appending

  APPEND => YES | <u>NO</u>

- Truncation of the name by DOS

  TRUNCATE => YES | <u>NO</u>

- DIRECT_IO on UNCONSTRAINED objects

  RECORD_SIZE => *size_in_bytes*

where:

> **BUFFER_SIZE:** Controls the size of the internal buffer. This option is not supported for DIRECT_IO. The default value is 1024. This option has no effect when used by TEXT_IO with an external file that is a character device, in which case the size of the buffer will be 0.

APPEND: If YES output is appended to the end of the existing file. If NO output overwrites the existing file. This option is not supported for DIRECT_IO. The default is NO.

TRUNCATE: If YES the file name will be automatically truncated if it is bigger than 8 characters. The default value is NO, meaning that the exception NAME_ERROR will be raised if the name is too long.

RECORD_SIZE: This option is supported only for DIRECT_IO. This attribute controls the logical record length of the external file.

- When DIRECT_IO is instantiated with an unconstrained type the user is required to specify the RECORD_SIZE attribute (otherwise USE_ERROR will be raised). The value given must be larger or equal to the largest record which is going to written. If a larger record is processed the exception USE_ERROR will be raised.

- When DIRECT_IO is instantiated with a constrained type the user is not required to specify the RECORD_SIZE but if the RECORD_SIZE is specified the only possible value would be the element size in bytes. Any other values will raise USE_ERROR.

The exception USE_ERROR is raised if the form STRING in not correct or if a non supported attribute for a given package is used.

Example:

```
FORM => "TRUNCATE => YES, APPEND => YES, BUFFER_SIZE => 20480"
```

8.4 Sequential Files

For sequential access the file is viewed as a sequence of values that are transferred in the order of their appearance (as produced by the program or run-time environment). This is sometimes called a *stream* file in other operating systems. Each object in a sequential file has the same binary representation as the Ada object in the executable program.

8.5 Direct Files

For direct access the file is viewed as a set of elements occupying consecutive positions in a linear order. The position of an element in a direct file is specified by its index, which is an integer of subtype POSITIVE_COUNT.

DIRECT_IO only allows input-output for constrained types. If DIRECT_IO is instantiated for an unconstrained type, all calls to CREATE or OPEN will raise USE_ERROR. Each object in a direct file will have the same binary representation as the Ada object in the executable program. All elements within the file will have the same length.

8.6 Text Files

All text file column numbers, line numbers, and .i.Page numbers;page numbers are values of the subtype .i.POSITIVE_COUNT;POSITIVE_COUNT.

Note that due to the definitions of line terminator, page terminator, and file terminator in the RM, and the method used to mark the end of file under DOS, some ASCII files do not represent well-formed TEXT_IO files.

A text file is buffered by the *Runtime Executive* unless

- it names a device (for example, CON or PRN).
- it is STANDARD_INPUT or STANDARD_OUTPUT band has not been redirected.

If not redirected, prompts written to STANDARD_OUTPUT with the procedure PUT will appear before (or when) a GET (or GET_LINE) occurs.

The functions END_OF_PAGE and END_OF_FILE always return FALSE when the file is a device, which includes the use of the file CON, and STANDARD_INPUT when it is not redirected. Programs which would like to check for end of file when the file may be a device should handle the exception END_ERROR instead, as in the following example:

Example

```
begin
    loop
        -- Display the prompt:
        TEXT_IO.PUT ("--> ");
        -- Read the next line:
        TEXT_IO.GET_LINE (COMMAND, LAST);
        -- Now do something with COMMAND (1 .. LAST)
    end loop;
exception
    when TEXT_IO.END_ERROR =>
        null;
end;
```

END_ERROR is raised for STANDARD_INPUT when ^Z (ASCII.SUB) is entered at the keyboard.

8.7 Access Protection of External Files

All DOS access protections exist when using files under DOS. If a file is open for read only access by one process it can not be opened by another process for read/write access.

8.8 The Need to Close a File Explicitly

The *Runtime Executive* will flush all buffers and close all open files when the program is terminated, either normally or through some exception.

However, the RM does not define what happens when a program terminates without closing all the opened files. Thus a program which depends on this feature of the *Runtime Executive* might have problems when ported to another system.

8.9 Limitation on the Procedure RESET

An internal file opened for input cannot be RESET for output. However, an internal file opened for output can be RESET for input, and can subsequently be RESET back to output.

8.10 Sharing of External Files and Tasking Issues

Several internal files can be associated with the same external file only if all the internal files are opened with mode IN_MODE. However, if a file is opened with mode OUT_MODE and then changed to IN_MODE with the RESET procedure, it cannot be shared.

Care should be taken when performing multiple input-output operations on an external file during tasking because the order of calls to the I/O primitives is unpredictable. For example, two strings output by TEXT_IO.PUT_LINE in two different tasks may appear in the output file with interleaved characters. Synchronization of I/O in cases such as this is the user's responsibility.

The TEXT_IO files STANDARD_INPUT and STANDARD_OUTPUT are shared by all tasks of an Ada program.

If TEXT_IO.STANDARD_INPUT is not redirected, it will not block a program on input. All tasks not waiting for input will continue running.

Section 9

Characteristics of Numeric Types

9.1 Integer Types

The ranges of values for integer types declared in package STANDARD are as follows:

SHORT_INTEGER	-128 .. 127	-- 2**7 - 1
INTEGER	-32768 .. 32767	-- 2**15 - 1
LONG_INTEGER	-2147483648 .. 2147483647	-- 2**31 - 1

For the packages DIRECT_IO and TEXT_IO, the range of values for types COUNT and POSITIVE_COUNT are as follows:

COUNT	0 .. 2147483647	-- 2**31 - 1
POSITIVE_COUNT	1 .. 2147483647	-- 2**31 - 1

For the package TEXT_IO, the range of values for the type FIELD is as follows:

FIELD	0 .. 255	-- 2**8 - 1

9.2 Floating Point Type Attributes

	FLOAT	LONG_FLOAT
DIGITS	6	15
MANTISSA	21	51
EMAX	84	204
EPSILON	9.53674E-07	8.88178E-16

LARGE	1.93428E+25	2.57110E+61
SAFE_EMAX	125	1021
SAFE_SMALL	1.17549E-38	2.22507E-308
SAFE_LARGE	4.25353E+37	2.24712E+307
FIRST	-3.40282E+38	-1.79769E+308
LAST	3.40282E+38	1.79769E+308
MACHINE_RADIX	2	2
MACHINE_EMAX	128	1024
MACHINE_EMIN	-125	-1021
MACHINE_ROUNDS	true	true
MACHINE_OVERFLOWS	false	false
SIZE	32	64

9.3 Attributes of Type DURATION

DURATION'DELTA	2.0 ** (-14)
DURATION'SMALL	2.0 ** (-14)
DURATION'FIRST	-131_072.0
DURATION'LAST	131_072.0
DURATION'LARGE	same as DURATION'LAST

Section 10

Other Implementation-Dependent Characteristics

10.1 Use of the Floating-Point Coprocessor

Floating point coprocessor instructions are used in programs that perform arithmetic on floating point values in some fixed point operations and when the FLOAT_IO or FIXED_IO packages of TEXT_IO are used. The mantissa of a fixed point value may be obtained through a conversion to an appropriate integer type. This conversion does not use floating point operations. Object code running on an 80286 using floating point instructions does not require the coprocessor, since software floating point emulation is provided. Object code running on an 80386 does require a coprocessor. Object code running on an 8086, 8088 or 80186 does require an 8087 coprocessor, since 8087 software emulation is not supported. See Appendix D of the *Application Developer's Guide* for more details.

The *Runtime Executive* will detect the absence of the floating point coprocessor if it is required by a program and will raise CONSTRAINT_ERROR.

10.2 Characteristics of the Heap

All objects created by allocators go into the heap. Also, portions of the *Runtime Executive* representation of task objects, including the task stacks, are allocated in the heap.

UNCHECKED_DEALLOCATION is implemented for all Ada access objects except access objects to tasks. Use of UNCHECKED_DEALLOCATION on a task object will lead to unpredictable results.

All objects whose visibility is linked to a subprogram, task body, or block have their storage reclaimed at exit, whether the exit is normal or due to an exception. Effectively pragma CONTROLLED is automatically applied to all access types. Moreover, all compiler temporaries on the heap (generated by such operations as function calls returning unconstrained arrays, or many concatenations) allocated in a scope are deallocated upon leaving the scope.

Note that the programmer may force heap reclamation of temporaries associated with any statements by enclosing the statement in a begin .. end block. This is especially useful when complex concatenations or other heap-intensive operations are performed in loops, and can reduce or eliminate STORAGE_ERRORs that might otherwise occur.

The maximum size of the heap is limited only by available memory. This includes the amount of physical memory (RAM) and the amount of virtual memory (hard disk swap space).

10.3 Characteristics of Tasks

The default task stack size is 1K bytes (32K bytes for the environment task), but by using the Binder option STACK.TASK the size for all task stacks in a program may be set to a size from 1K bytes to 64K bytes.

Normal priority rules are followed for preemption, where PRIORITY values are in the range 1 .. 10. A task with *undefined* priority (no pragma PRIORITY) is considered to be lower than priority 1.

The maximum number of active tasks is restricted only by memory usage.

The accepter of a rendezvous executes the accept body code in its own stack. Rendezvous with an empty accept body (for synchronization) does not cause a context switch.

The main program waits for completion of all tasks dependent upon library packages before terminating.

Abnormal completion of an aborted task takes place immediately, except when the abnormal task is the caller of an entry that is engaged in a rendezvous, or if it is in the process of activating some tasks. Any such task becomes abnormally completed as soon as the state in question is exited.

The message

```
GLOBAL BLOCKING SITUATION DETECTED
```

is printed to STANDARD_OUTPUT when the *Runtime Executive* detects that no further progress is possible for any task in the program. The execution of the program is then abandoned.

10.4 Definition of a Main Subprogram

A library unit can be used as a main subprogram if and only if it is a procedure that is not generic and that has no formal parameters.

The Compiler imposes no additional ordering constraints on compilations beyond those required by the language.

Section 11

Limitations

11.1 Compiler Limitations

- The maximum identifier length is 255 characters.
- The maximum line length is 255 characters.
- The maximum number of unique identifiers per compilation unit is 2500.
- The maximum number of compilation units in a library is 1000.
- The maximum number of Ada libraries in a family is 15.

11.2 Hardware Related Limitations

- The maximum amount of data in the heap is limited only by available memory.
- If an unconstrained record type can exceed 4096 bytes, the type is not permitted (unless constrained) as the element type in the definition of an array or record type.
- The maximum size of the generated code for a single compilation unit is 65535 bytes.
- The maximum size of a single array or record object is 65522 bytes. An object bigger than 4096 bytes will be indirectly allocated. Refer to ALLOCATION parameter in the COMPILE command. (Section 4.2 of the *User's Guide.*)
- The maximum size of a single stack frame is 32766 bytes, including the data for inner package subunits unnested to the parent frame.
- The maximum amount of data in the global data area is 65535 bytes, including compiler generated data that goes into the GDA (about 8 bytes per compilation unit plus 4 bytes per externally visible subprogram).

INDEX